To Err is Human:

The Complexity of Security Failures

Kelly Shortridge (@swagitda_)

Hacktivity 2019 Keynote

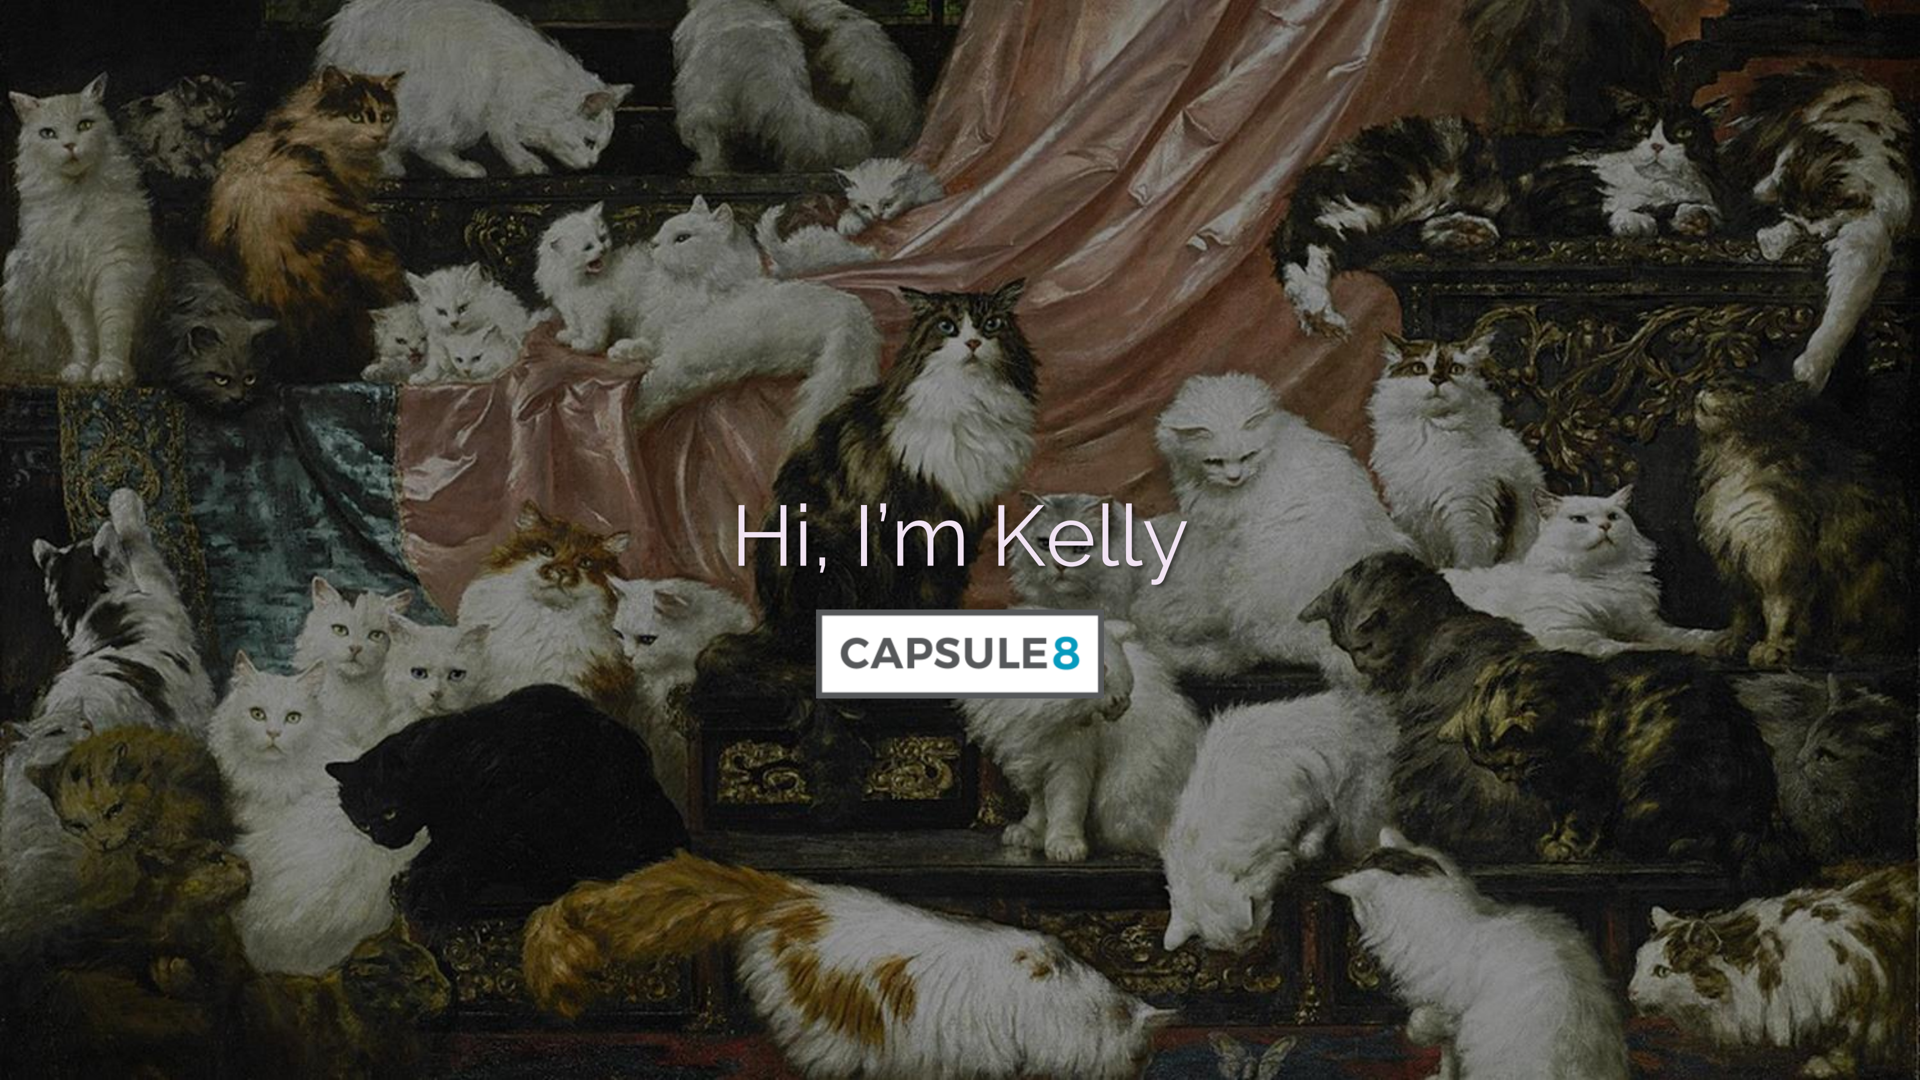

Hi, I’m Kelly

CAPSULE8

“To err is human; to forgive, divine.”
– Alexander Pope

Humans make mistakes. It’s part of our nature (it’s mostly a feature, not a bug)

Infosec's mistake: operating as if you can force humans to never err

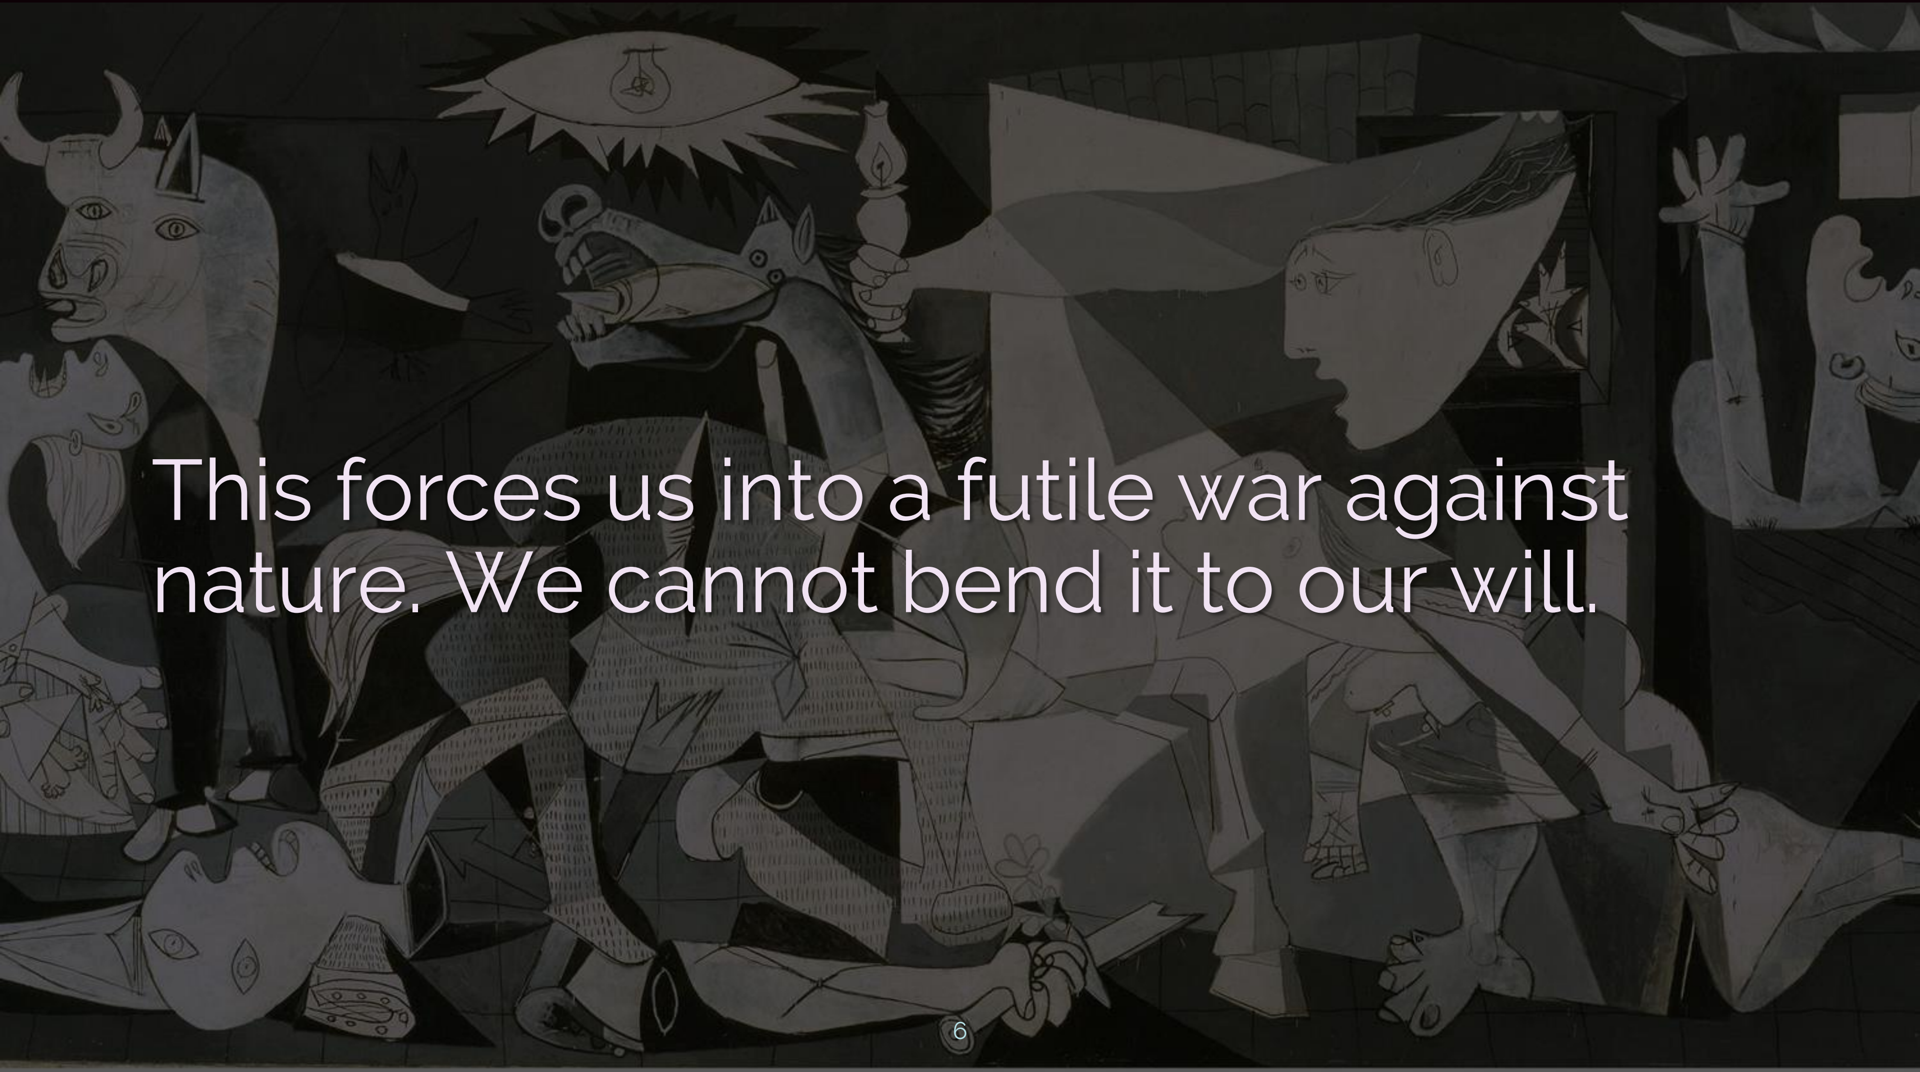

This forces us into a futile war against nature. We cannot bend it to our will.

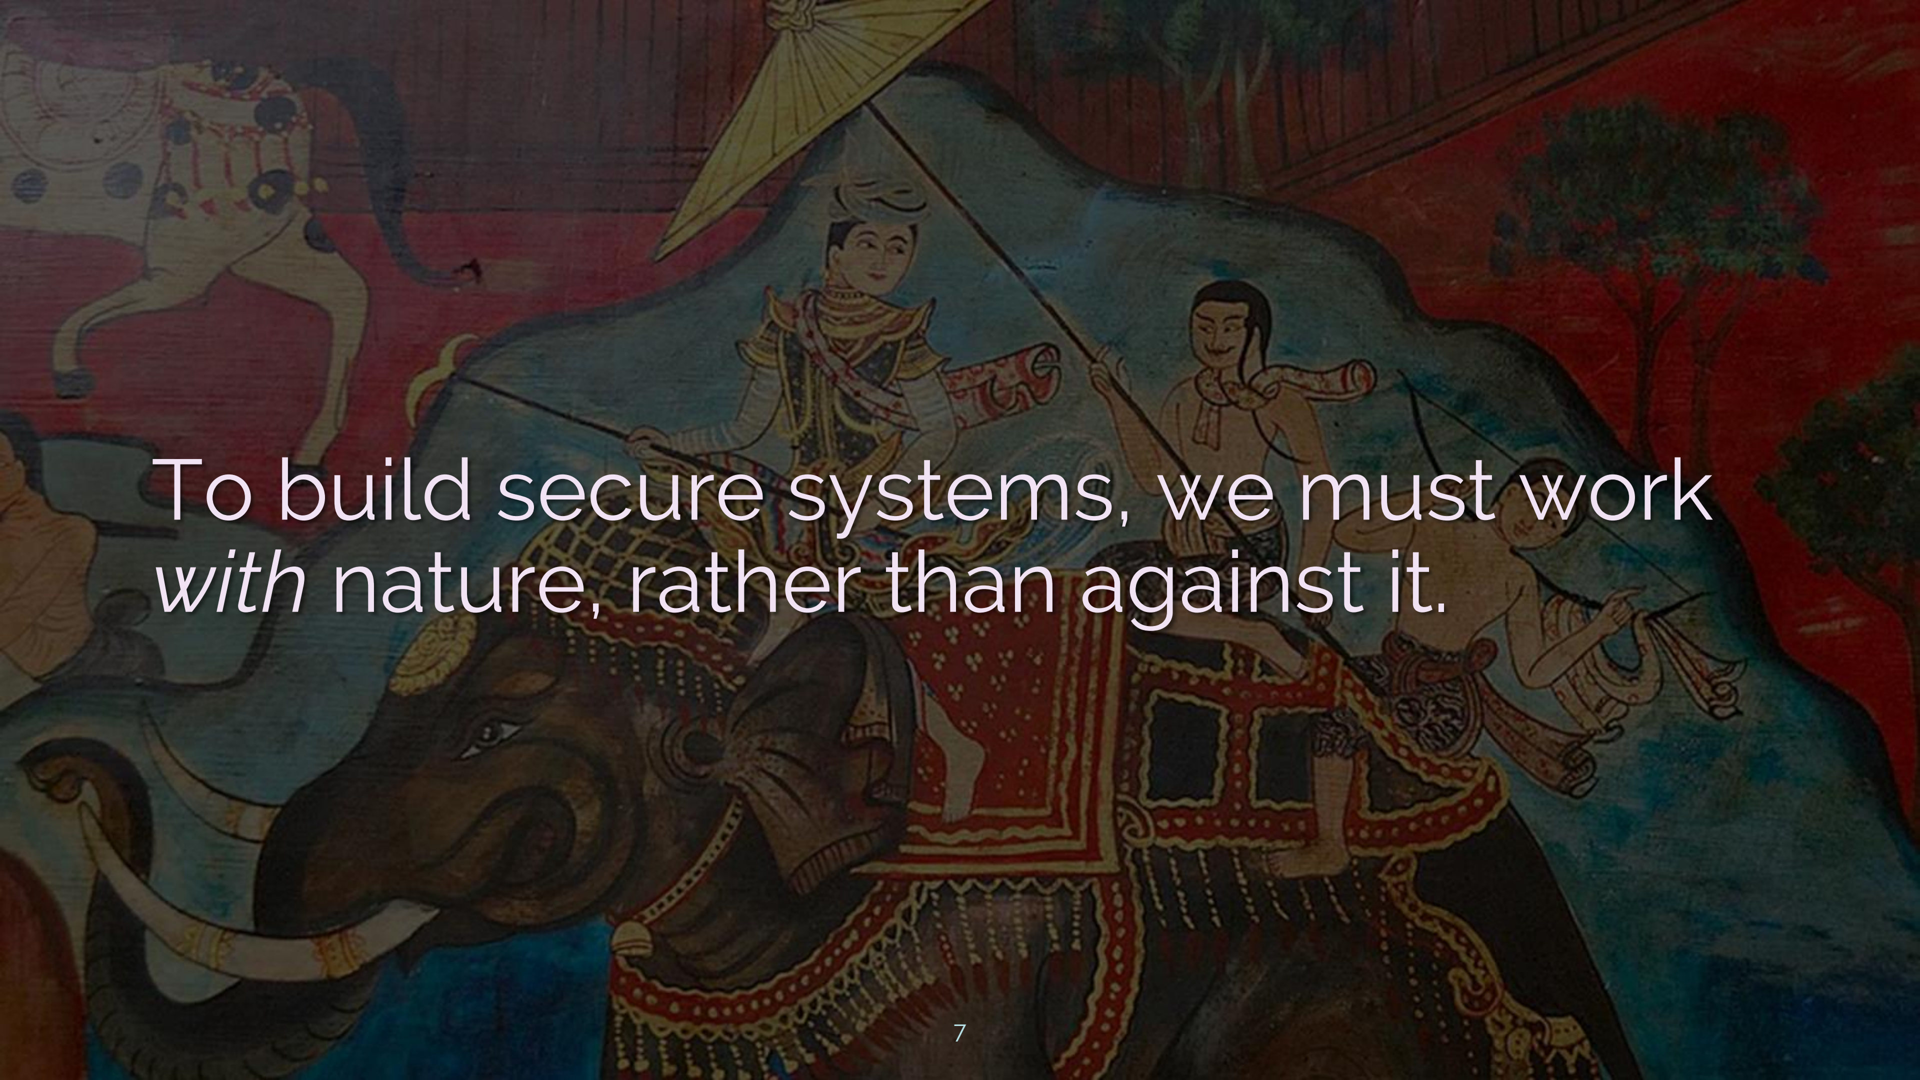

To build secure systems, we must work *with* nature, rather than against it.

1. Clearing the Err
2. Hindsight & Outcome Bias
3. Unhealthy Coping Mechanisms
4. Making Failure Epic

Clearing the Err

Error: an action that leads to failure or that deviates from expected behavior

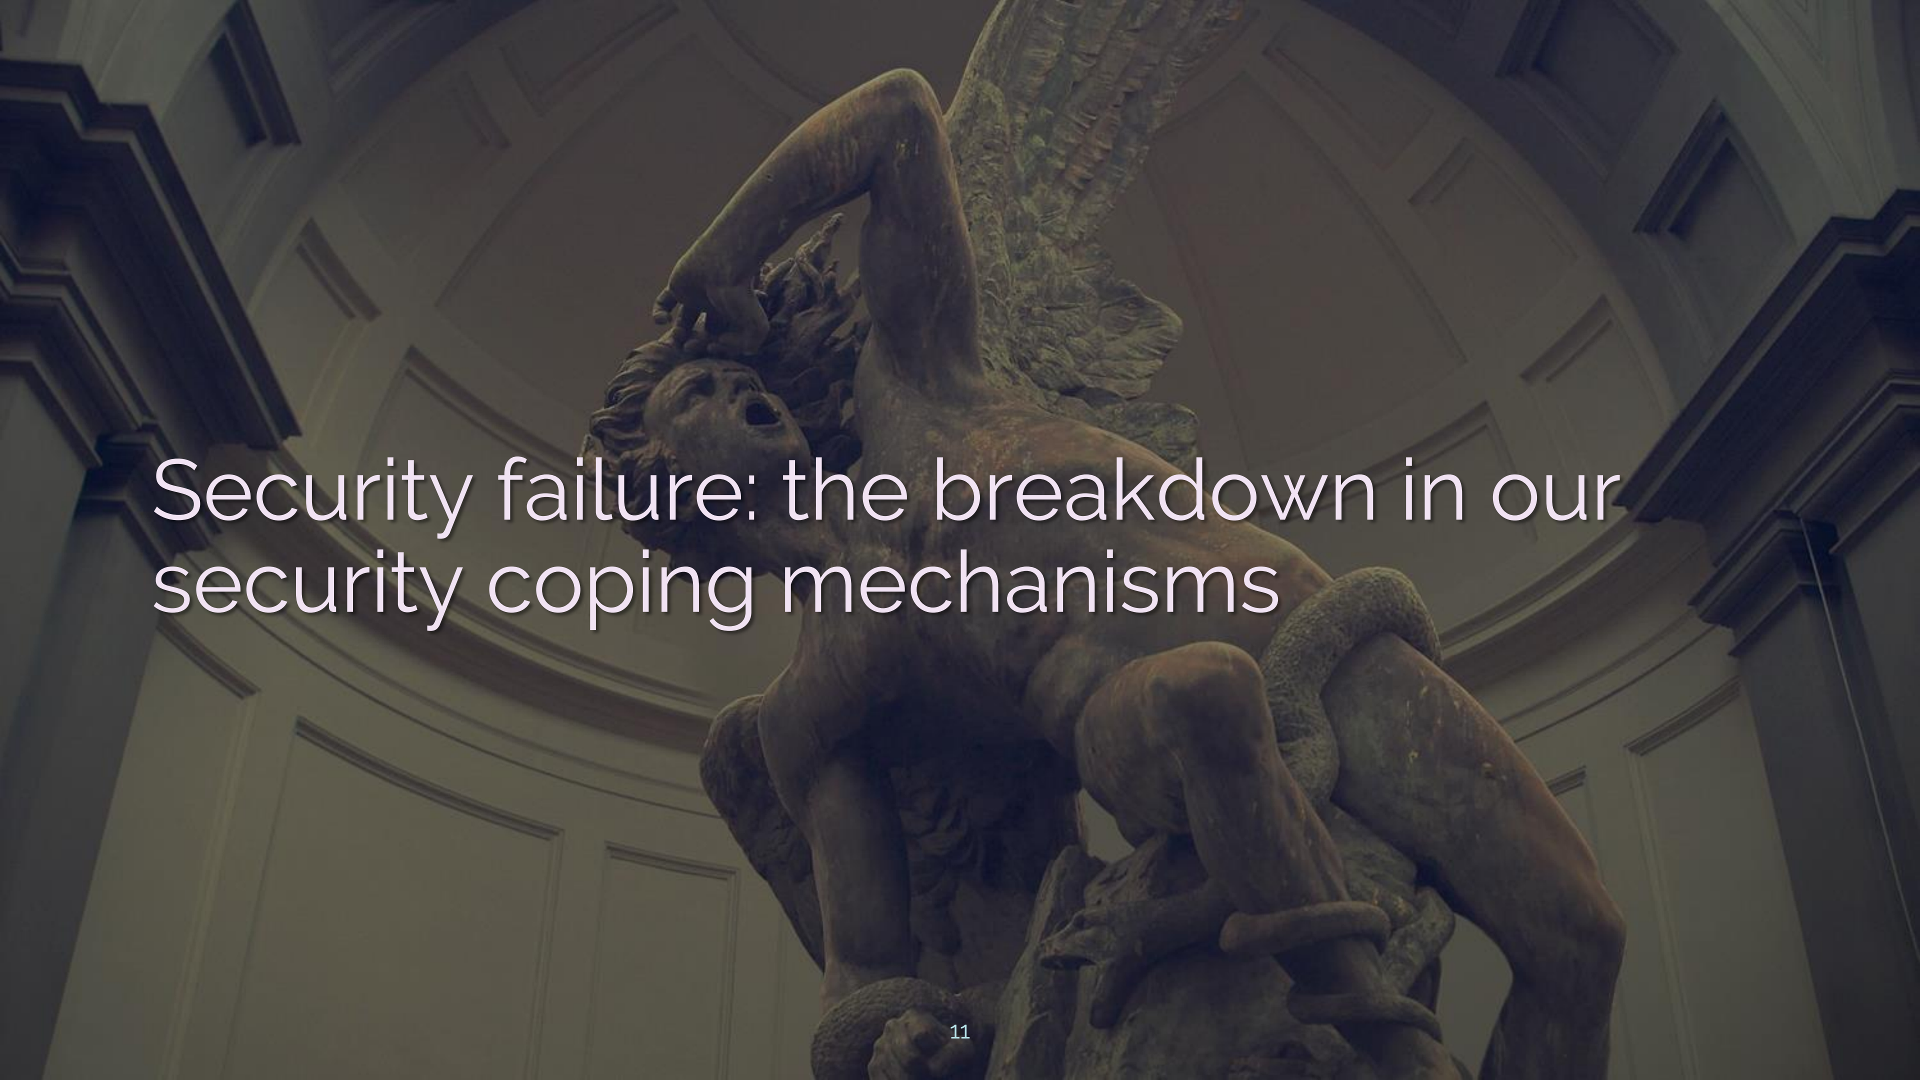

Security failure: the breakdown in our security coping mechanisms

“Human error” involves subjective expectations, including in infosec

Understanding why incidents happened is essential, but blame doesn't help

Aviation, manufacturing, & healthcare are already undergoing this revolution

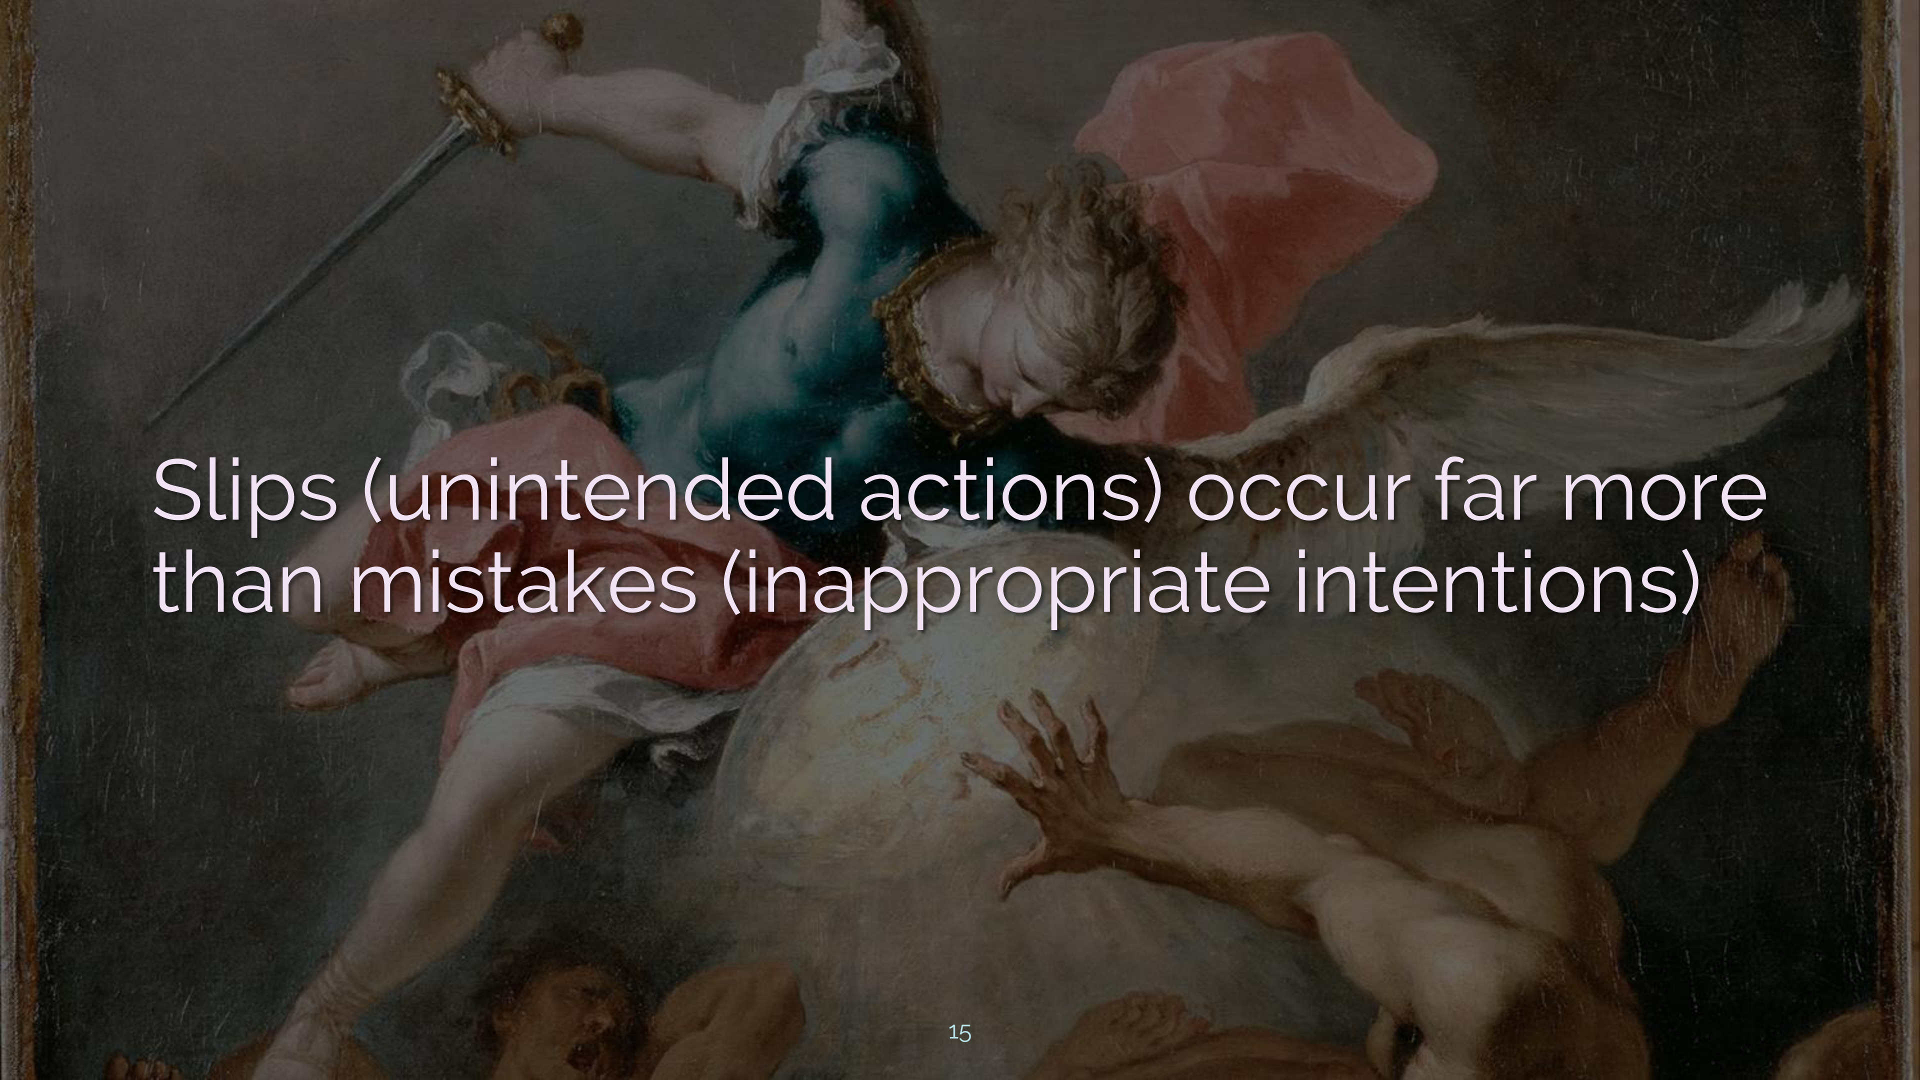

Slips (unintended actions) occur far more than mistakes (inappropriate intentions)

The term “human error” is less grounded to reality than we believe…

Hindsight & Outcome Bias

Cognitive biases represent mental shortcuts that are optimal for evolution

We learn from the past to progress, but our “lizard brain” can take things too far

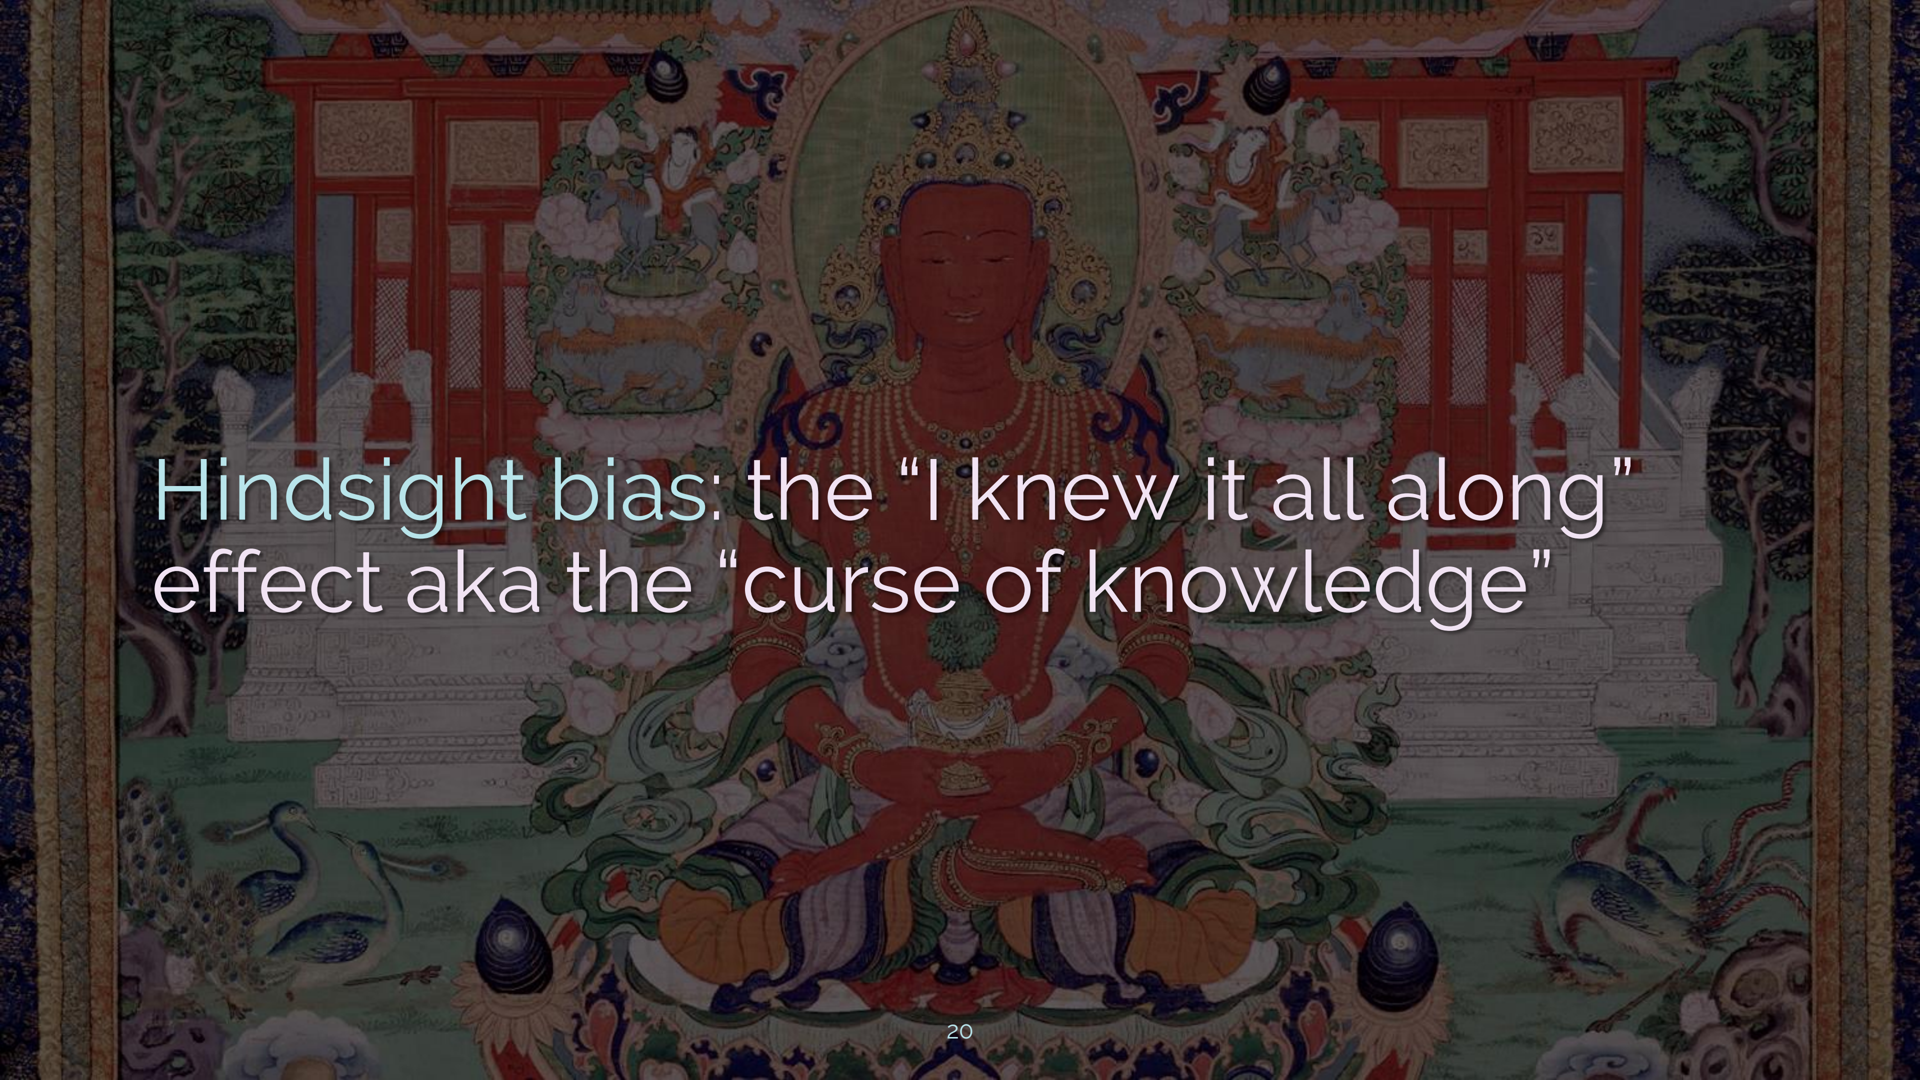

Hindsight bias: the “I knew it all along” effect aka the “curse of knowledge”

People overestimate their predictive abilities when lacking future knowledge e.g. skepticism of N.K. attribution for the Sony Pictures leak; now it is “obvious”

Outcome bias: judging a decision based on its eventual outcome

Instead, evaluate decisions based on what was known at that time

All decisions involve some level of risk.
Outcomes are largely based on chance.

We unfairly hold people accountable for events beyond their control

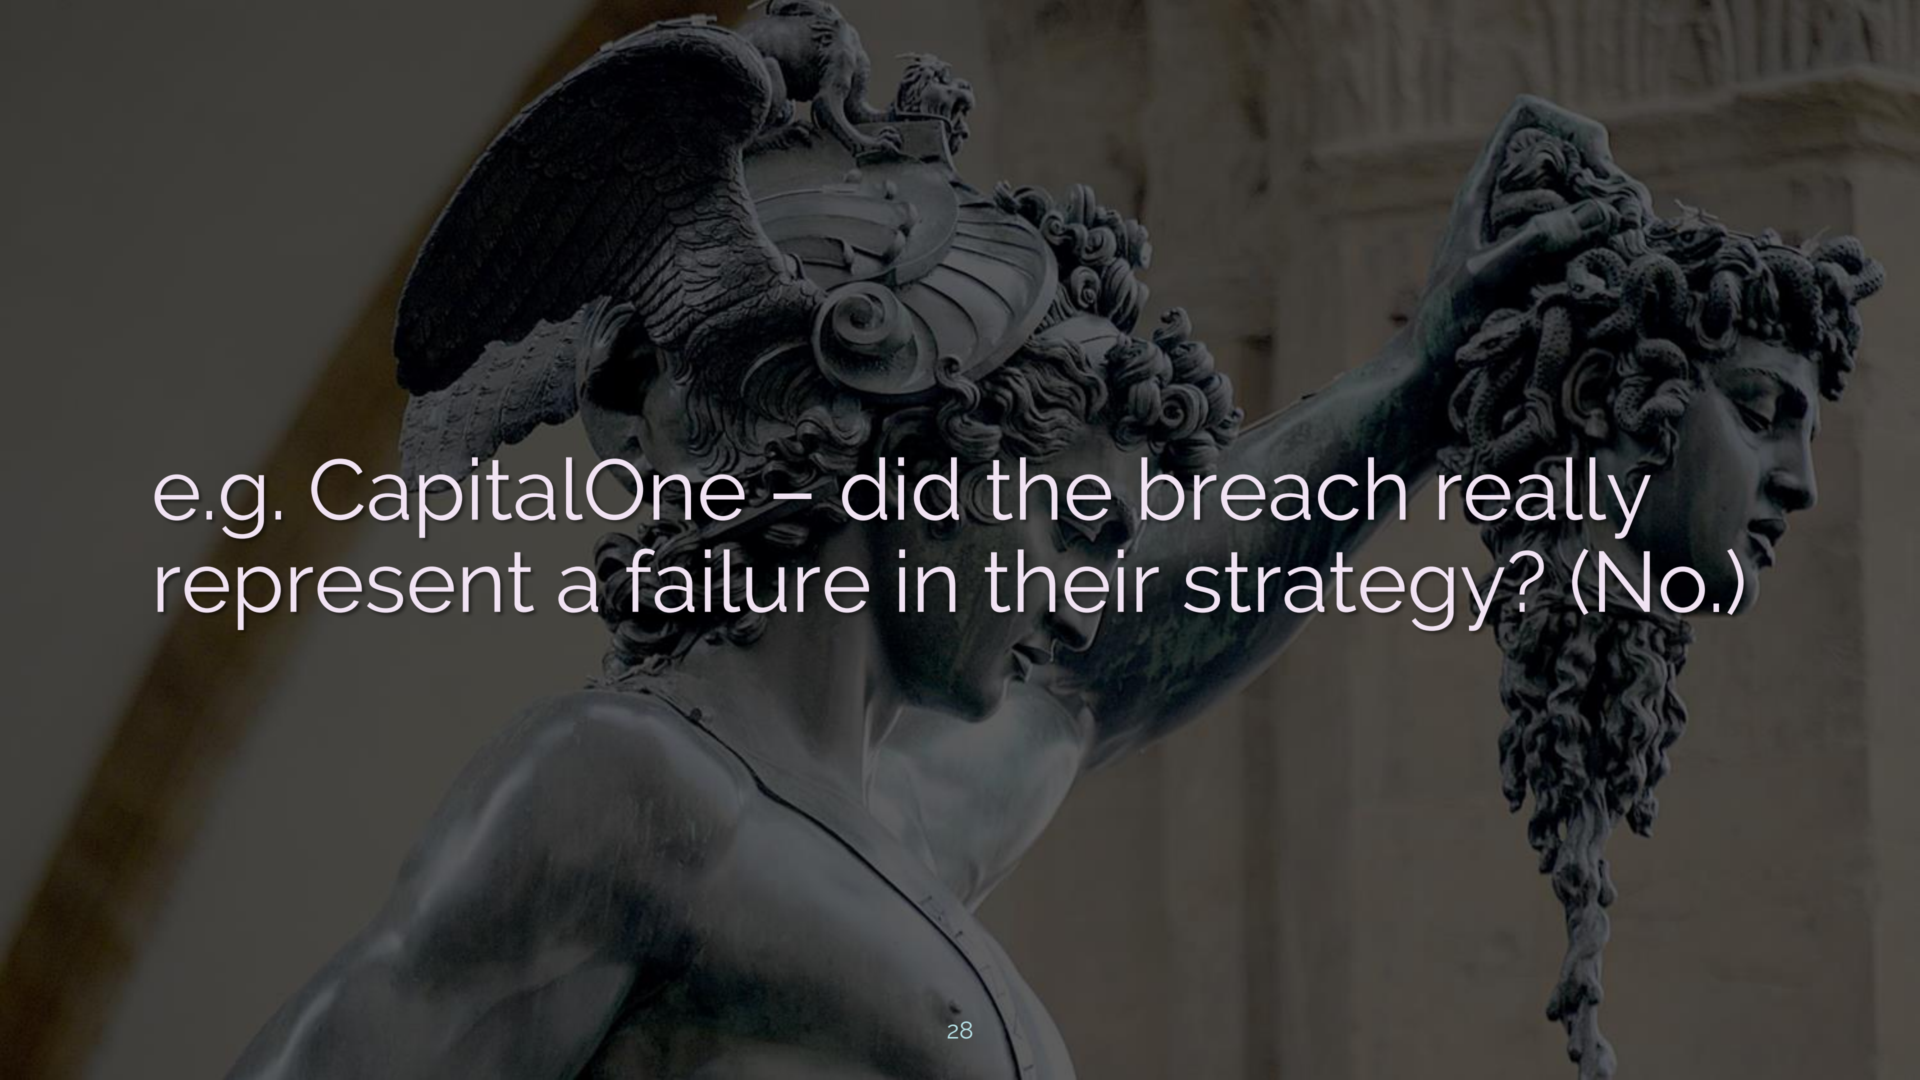

e.g. CapitalOne – did the breach really represent a failure in their strategy? (No.)

These biases change how we cope with failure…

Unhealthy Coping Mechanisms

Unhealthy coping mechanism #1: Blaming “human error”

Infosec’s fav hobbies: PICNIC & PEBKAC

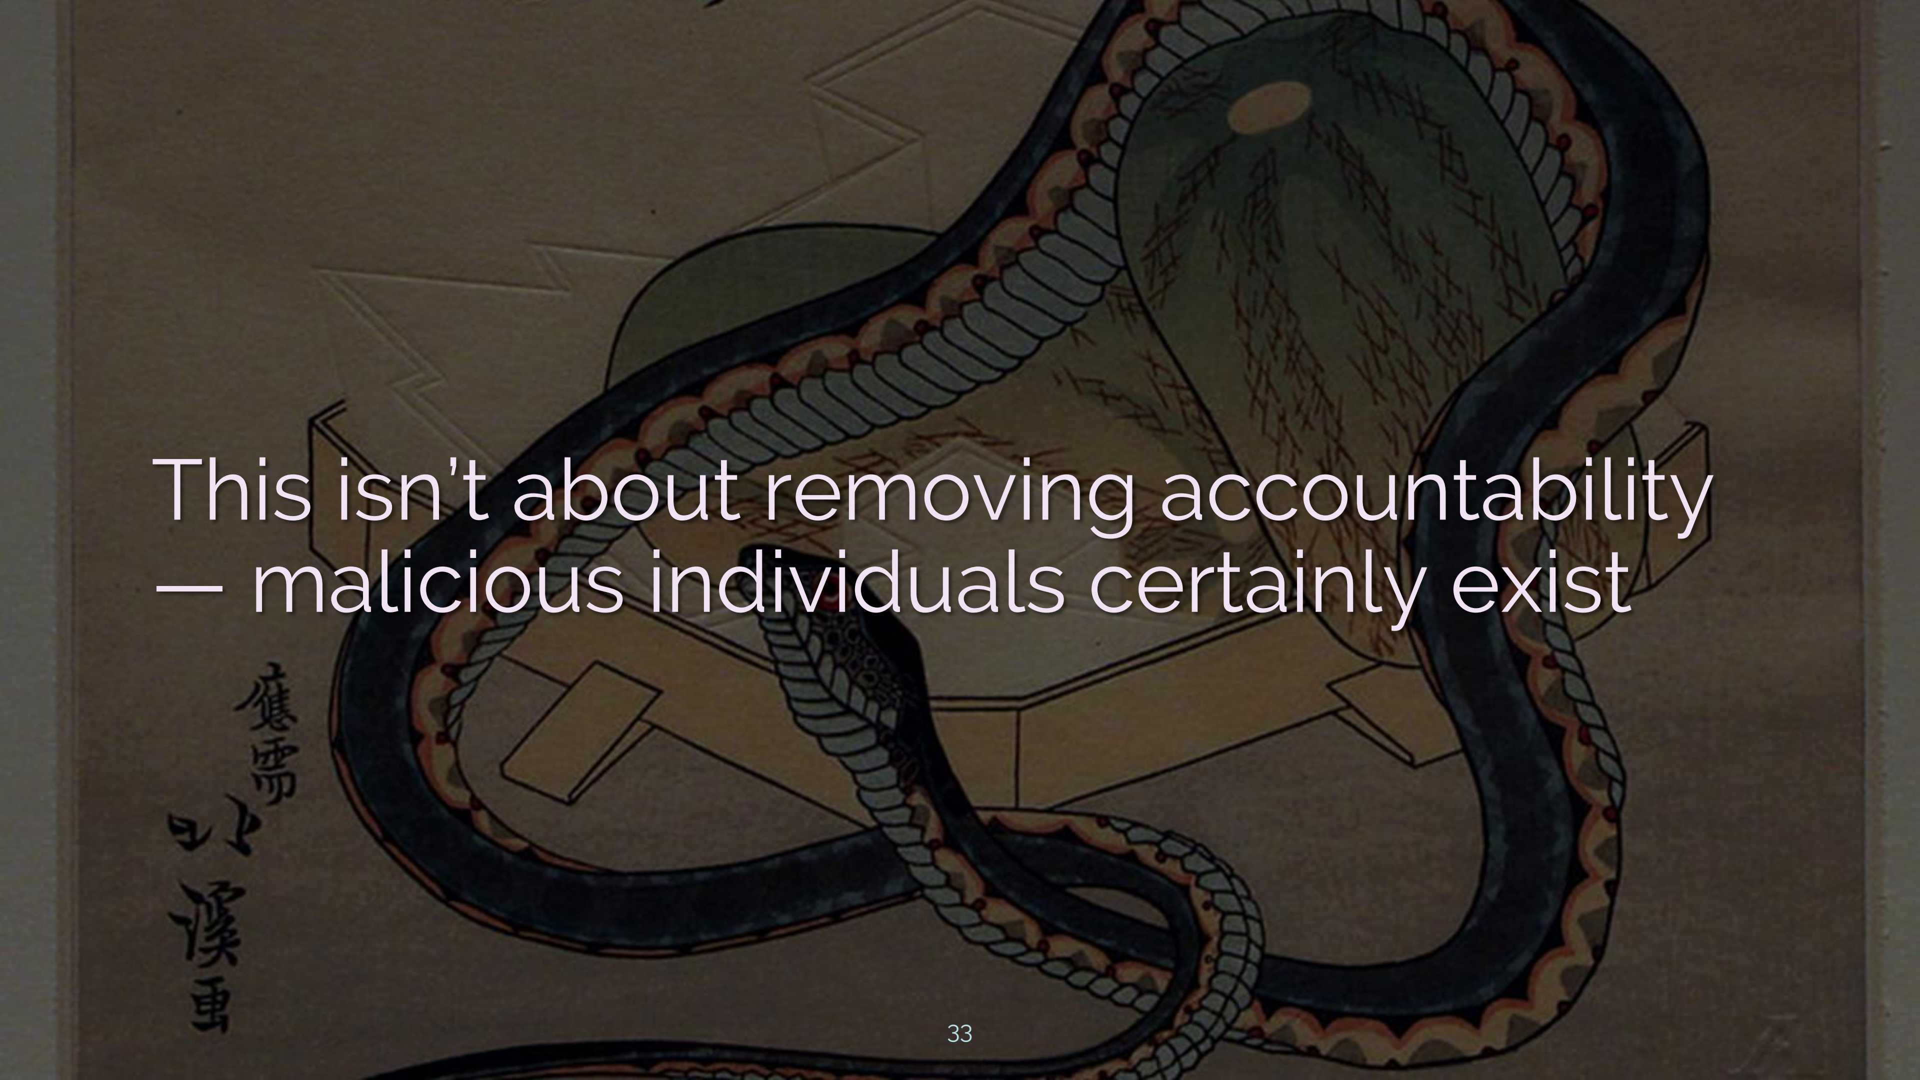

This isn’t about removing accountability — malicious individuals certainly exist

Fundamental attribution error: your actions reflect innate traits, mine don’t

"You are inattentive, sloppy, & naïve for clicking a link. I was just super busy."

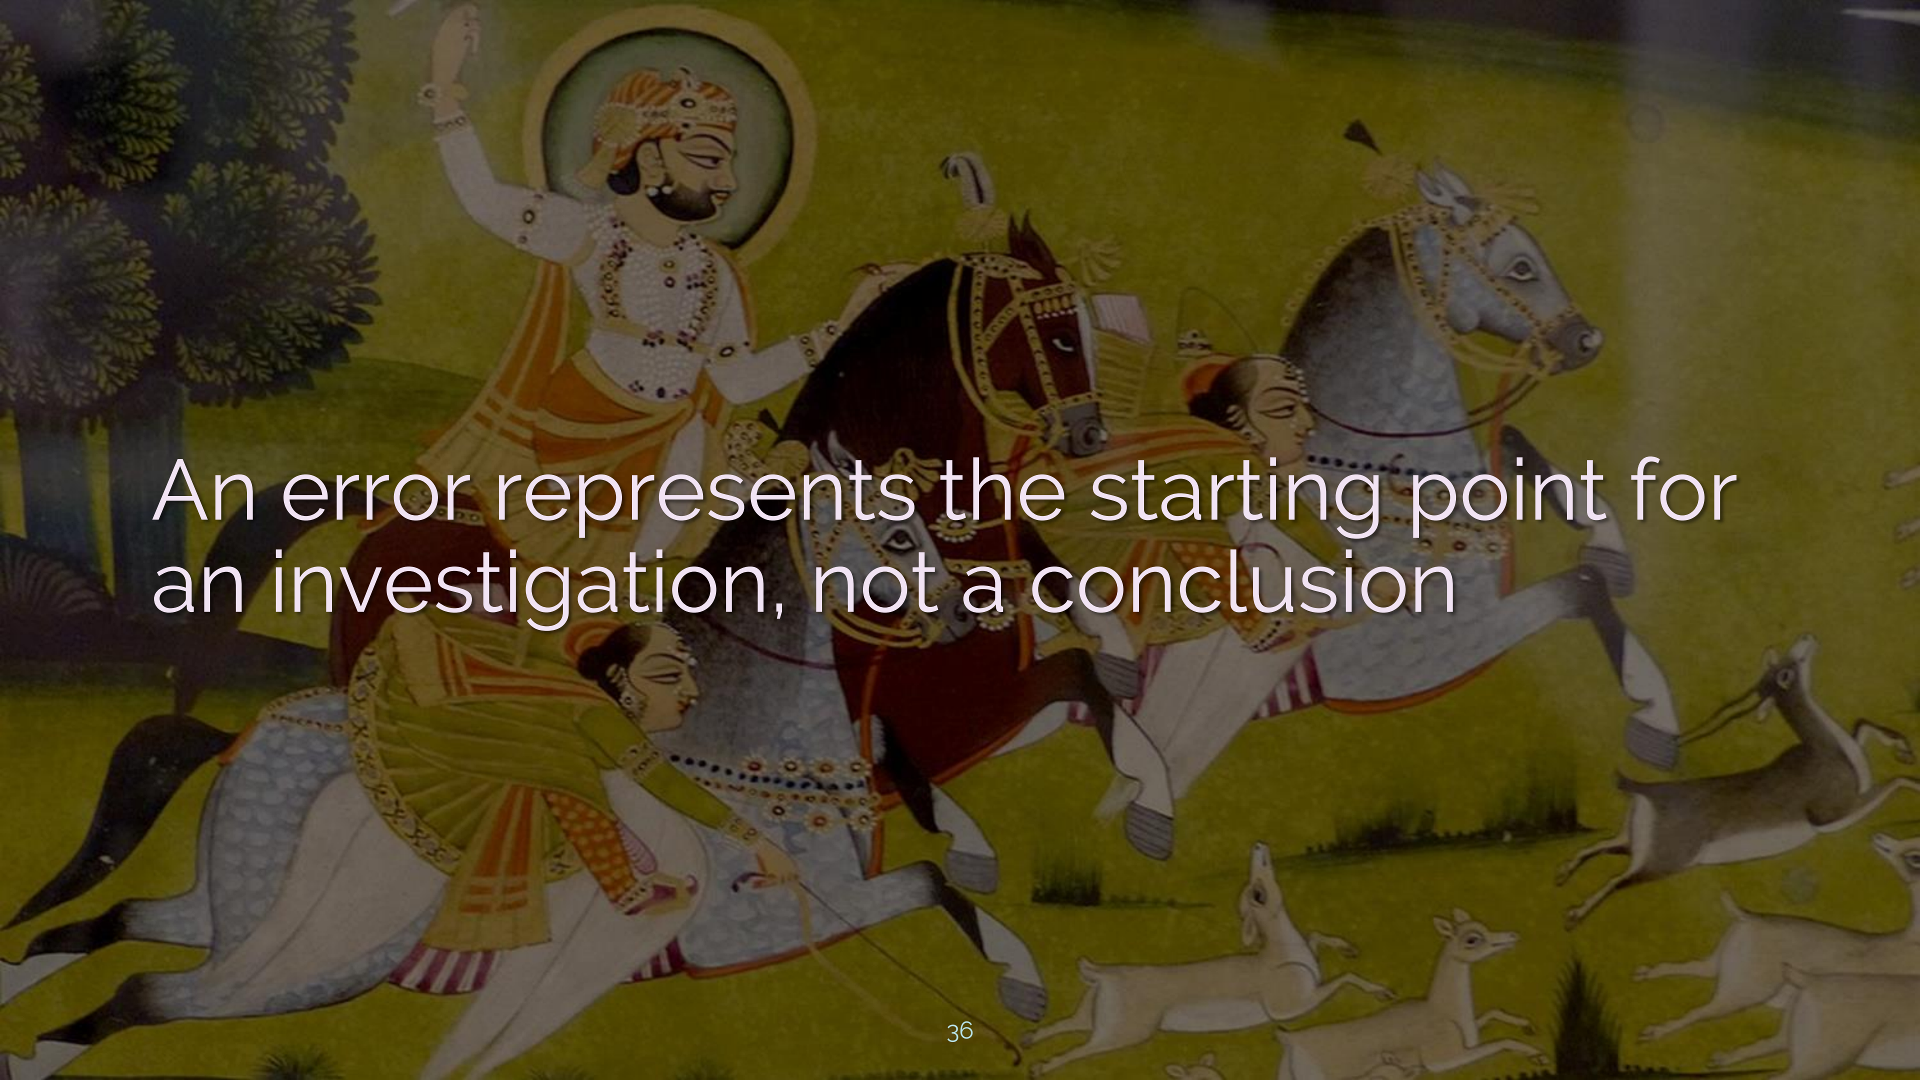

An error represents the starting point for an investigation, not a conclusion

“Why did they click the link?”
“Why did clicking a link lead to pwnage?”

These questions go unanswered if we accept the “human error” explanation

e.g. training devs to “care about security” completely misses the underlying issue

A “5 Whys” approach is a healthy start

Equifax’s ex-CEO blamed “human error” for the breach. He was wrong.

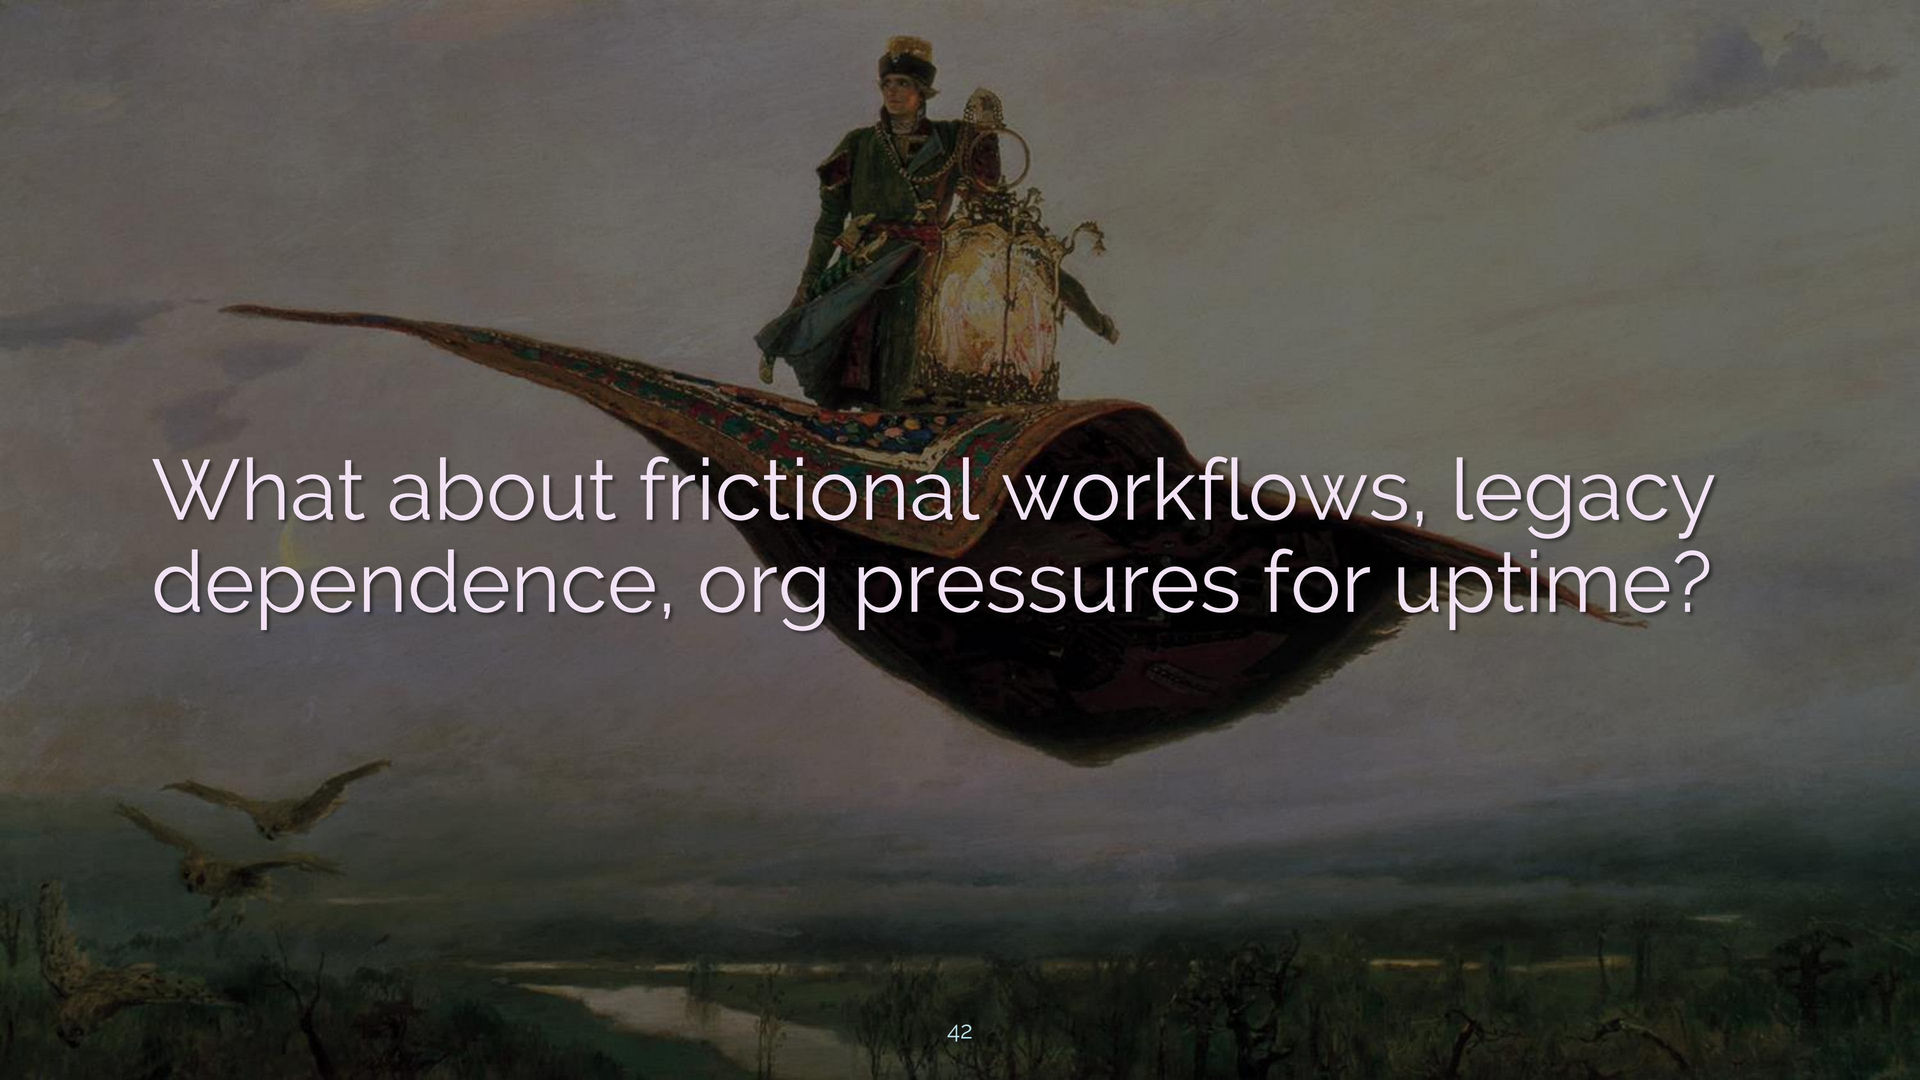

What about frictional workflows, legacy dependence, org pressures for uptime?

90% of breaches cite “human error” as the cause. That stat is basically useless.

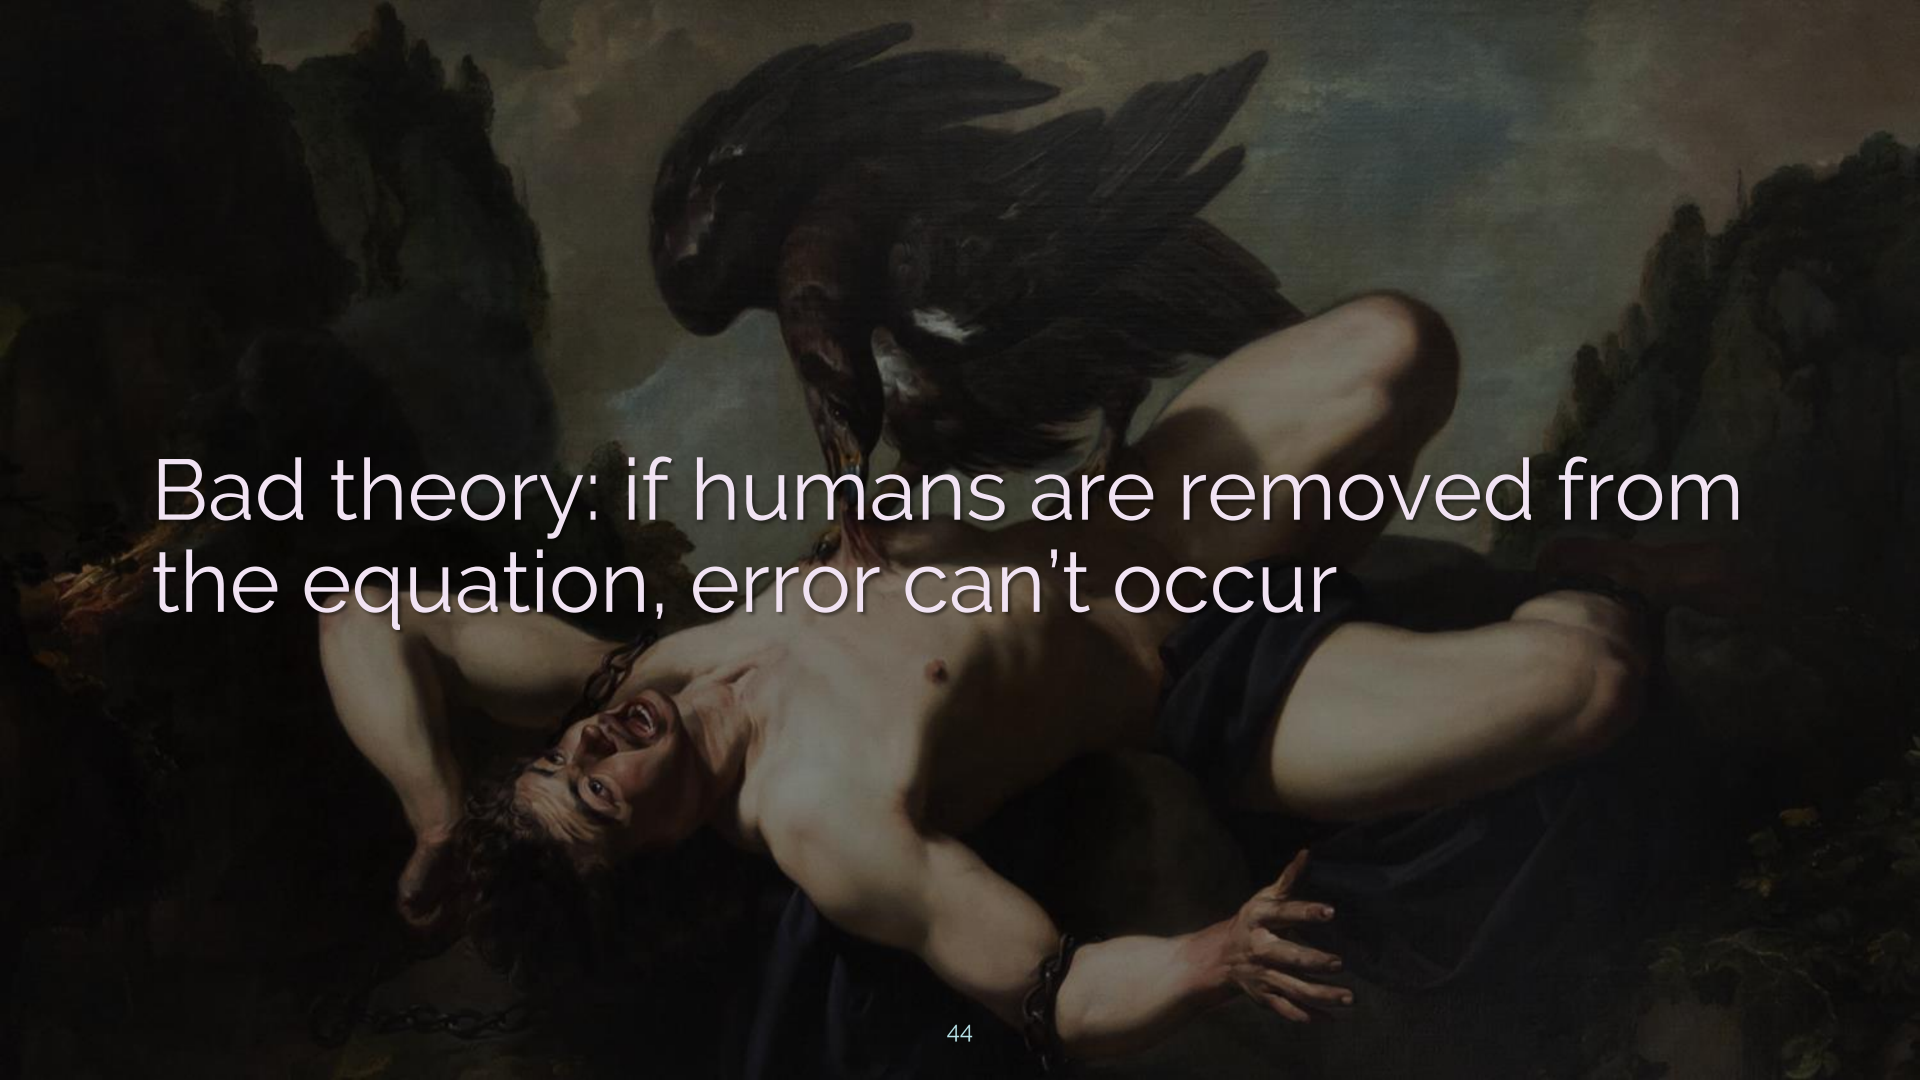

Bad theory: if humans are removed from the equation, error can’t occur

Unhealthy coping mechanism #2: Behavioral control

"An approach aimed at the individual is the equivalent of swatting individual mosquitoes rather than draining the swamp to address the source of the problem."

– Henriksen, et al.

“Policy violation” is a sneaky way to still rely on “human error” as an answer

The cornucopia of security awareness hullabaloo is a direct result of this

Solely restricting human behavior will never improve security outcomes.

We focus on forcing humans to fit our ideal mold vs. re-designing our systems

Formal policies are rarely written by those in the flow of work being policed

Infosec is mostly at the “blunt” end of systems; operators are at the “sharp” end

People tend to blame whomever resides closest to the error

Operator actions “add a final garnish to a lethal brew whose ingredients have already been long in the cooking.”

– James Reason

e.g. Equifax’s 48-hour patching policy that was very obviously not followed

Creating words on a piece of paper & expecting results is… ambitious

Discipline doesn't actually fix the "policy violation" cause (but it does scapegoat)

Case study: SS&C & BEC

Solely implementing controls to regulate human behavior doesn’t beget resilience

Post-WWII analysis: Improved design of cockpit controls won over pilot training

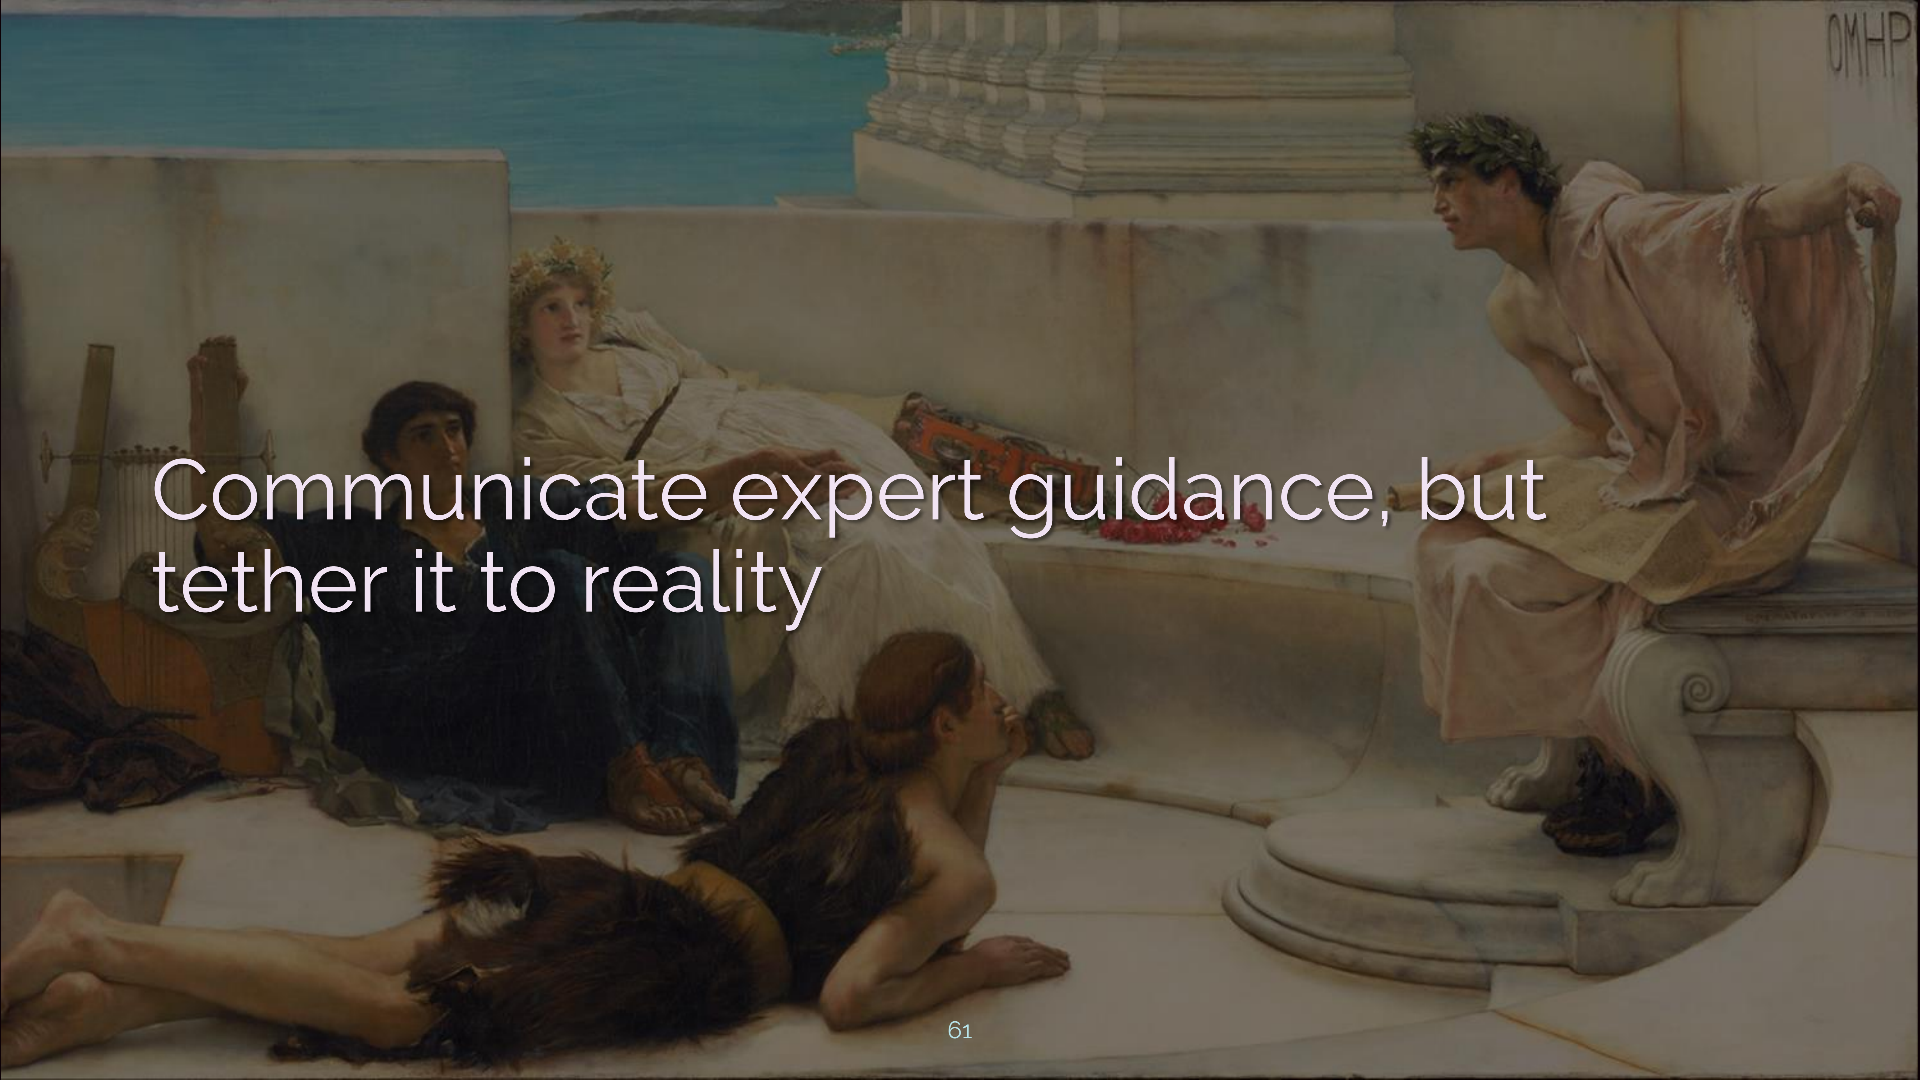

Communicate expert guidance, but tether it to reality

Checklists can be valuable aids *if* they're based on knowledge of real workflows

Policies must encourage safer contexts, not lord over behavior with an iron fist.

Unhealthy coping mechanism #3:
The just-world hypothesis

Attempting to find the ultimate causal seed of failure helps us cope with fear

The just world hypothesis: humans like believing the world is orderly & fair

The fact that the same things can lead to both success & failure isn’t a “just world”

Case Study: The Chernobyl disaster

Errors are really symptoms of pursuing goals while under resource constraints

How can security teams more productively deal with security failures?

Making Failure Epic

Infosec will progress when we ensure the easy way is the secure way

1. System perspective
2. Security UX
3. Chaos security engineering
4. Blameless culture

System perspective

Security failure is never the result of one factor, one vuln, or one dismissed alert

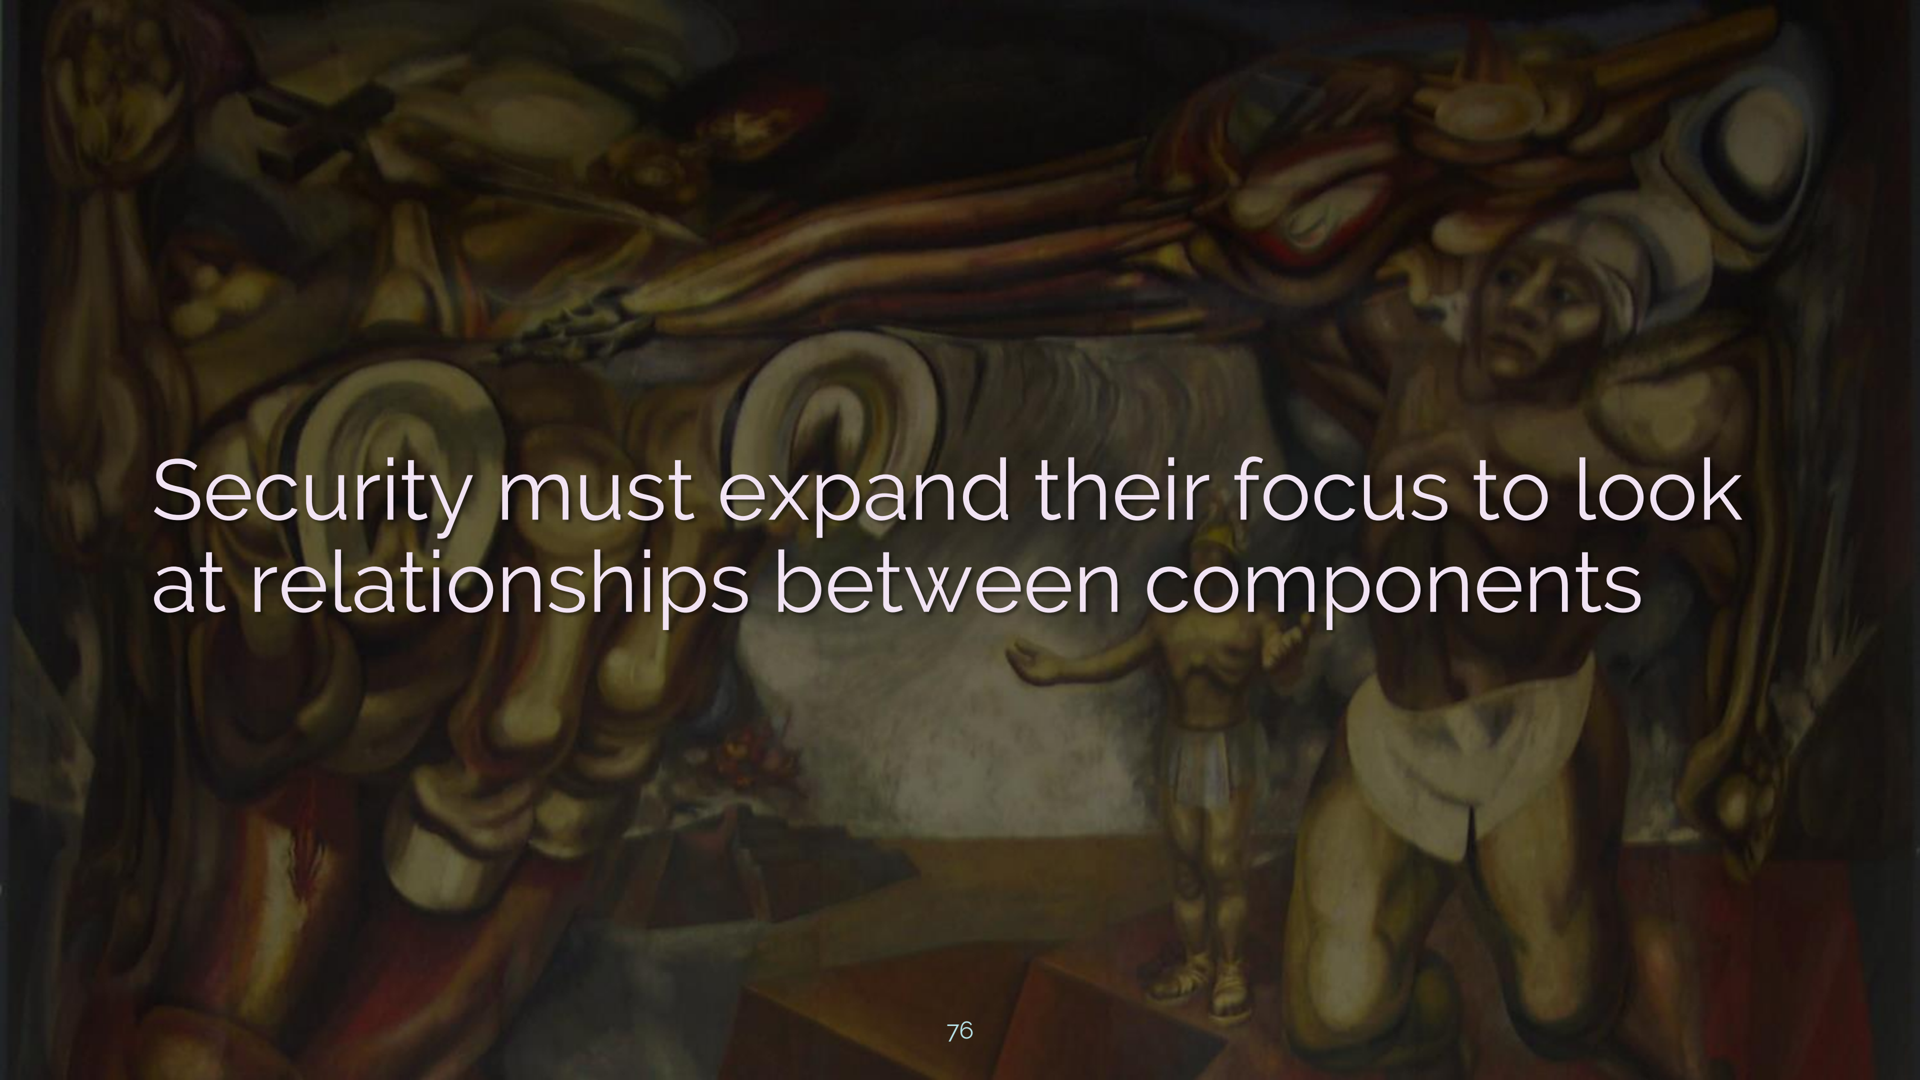

Security must expand their focus to look at relationships between components

A system is “a set of interdependent components interacting to achieve a common specified goal.”

“A narrow focus on operator actions, physical component failures, and technology may lead to ignoring some of the most important factors in terms of preventing future accidents”

– Nancy Leveson

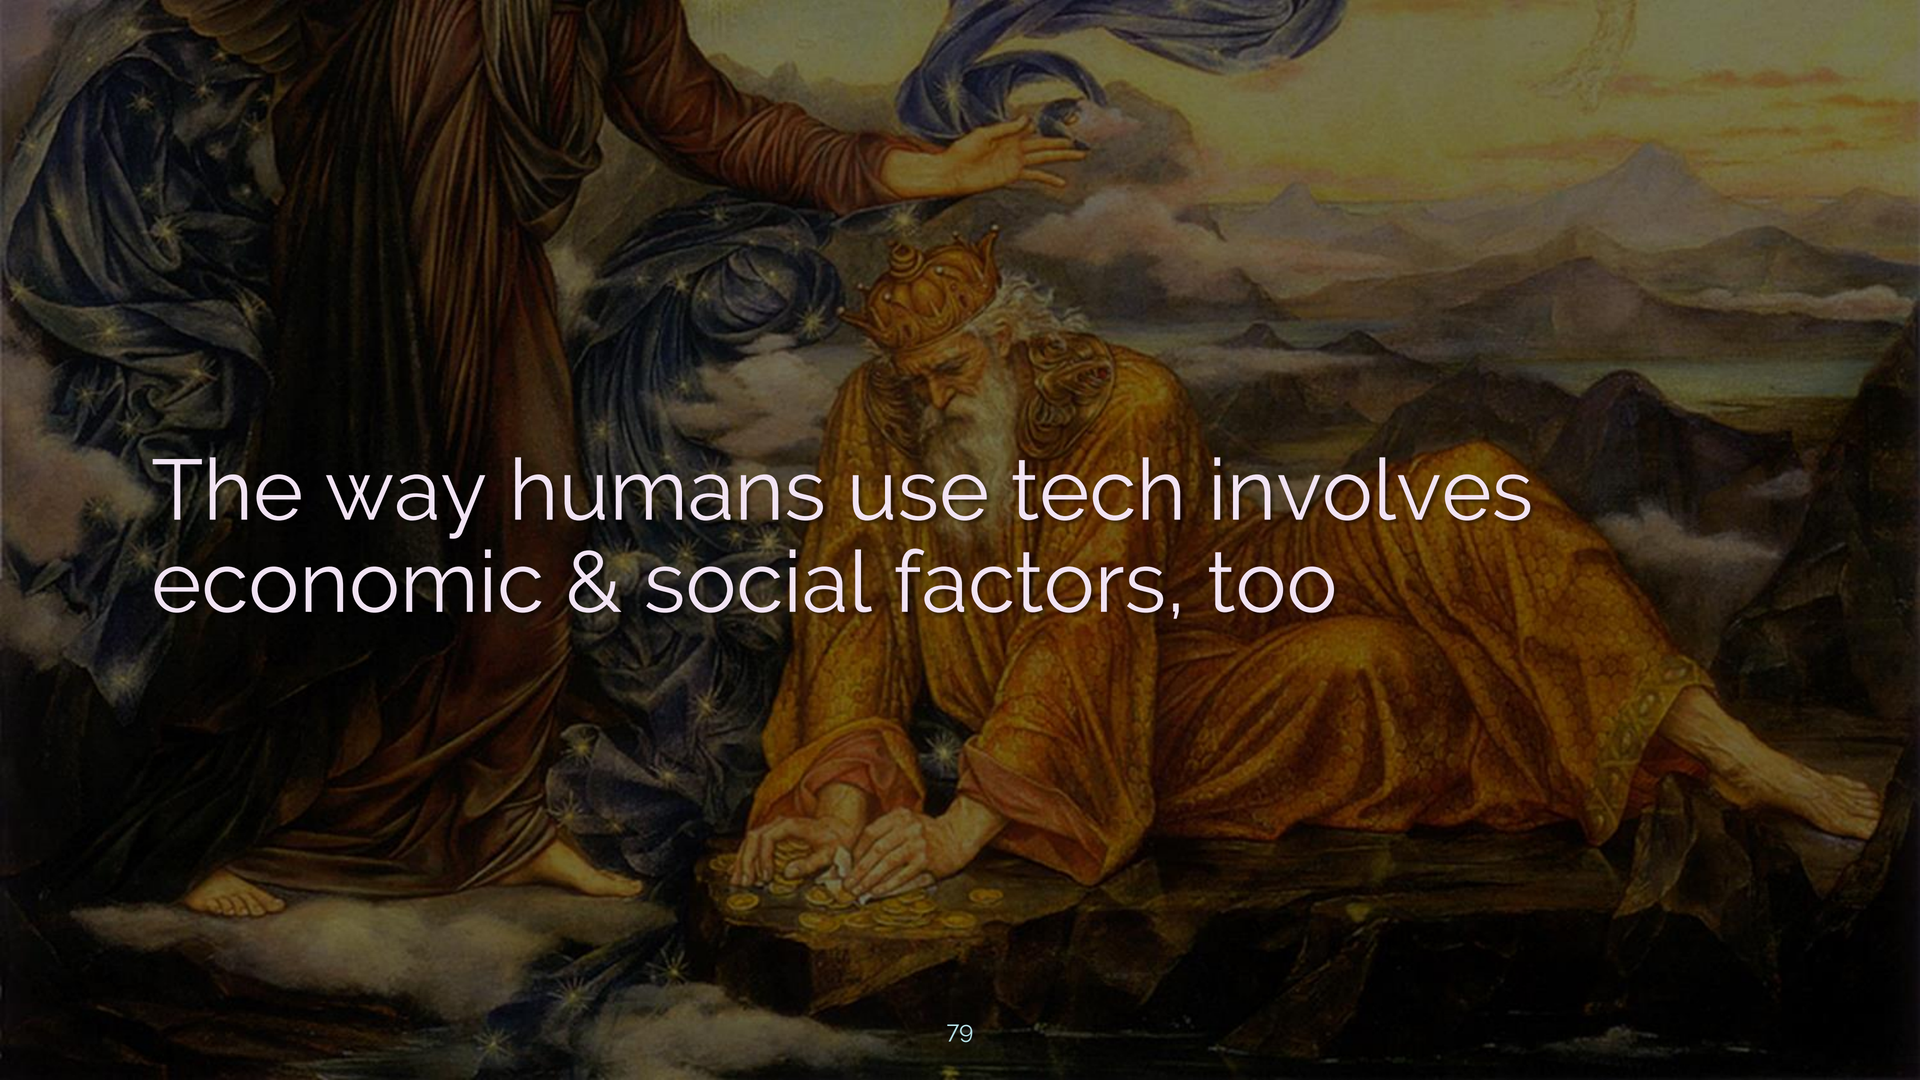

The way humans use tech involves economic & social factors, too

Economic factors: revenue & profit goals, compensation schemes, budgeting, etc.

Social factors: KPIs, expectations, what behavior is rewarded or punished, etc.

Pressure to do more work, faster is a vulnerability. So is a political culture.

Non-software vulns don’t appear in our threat models, but also erode resilience

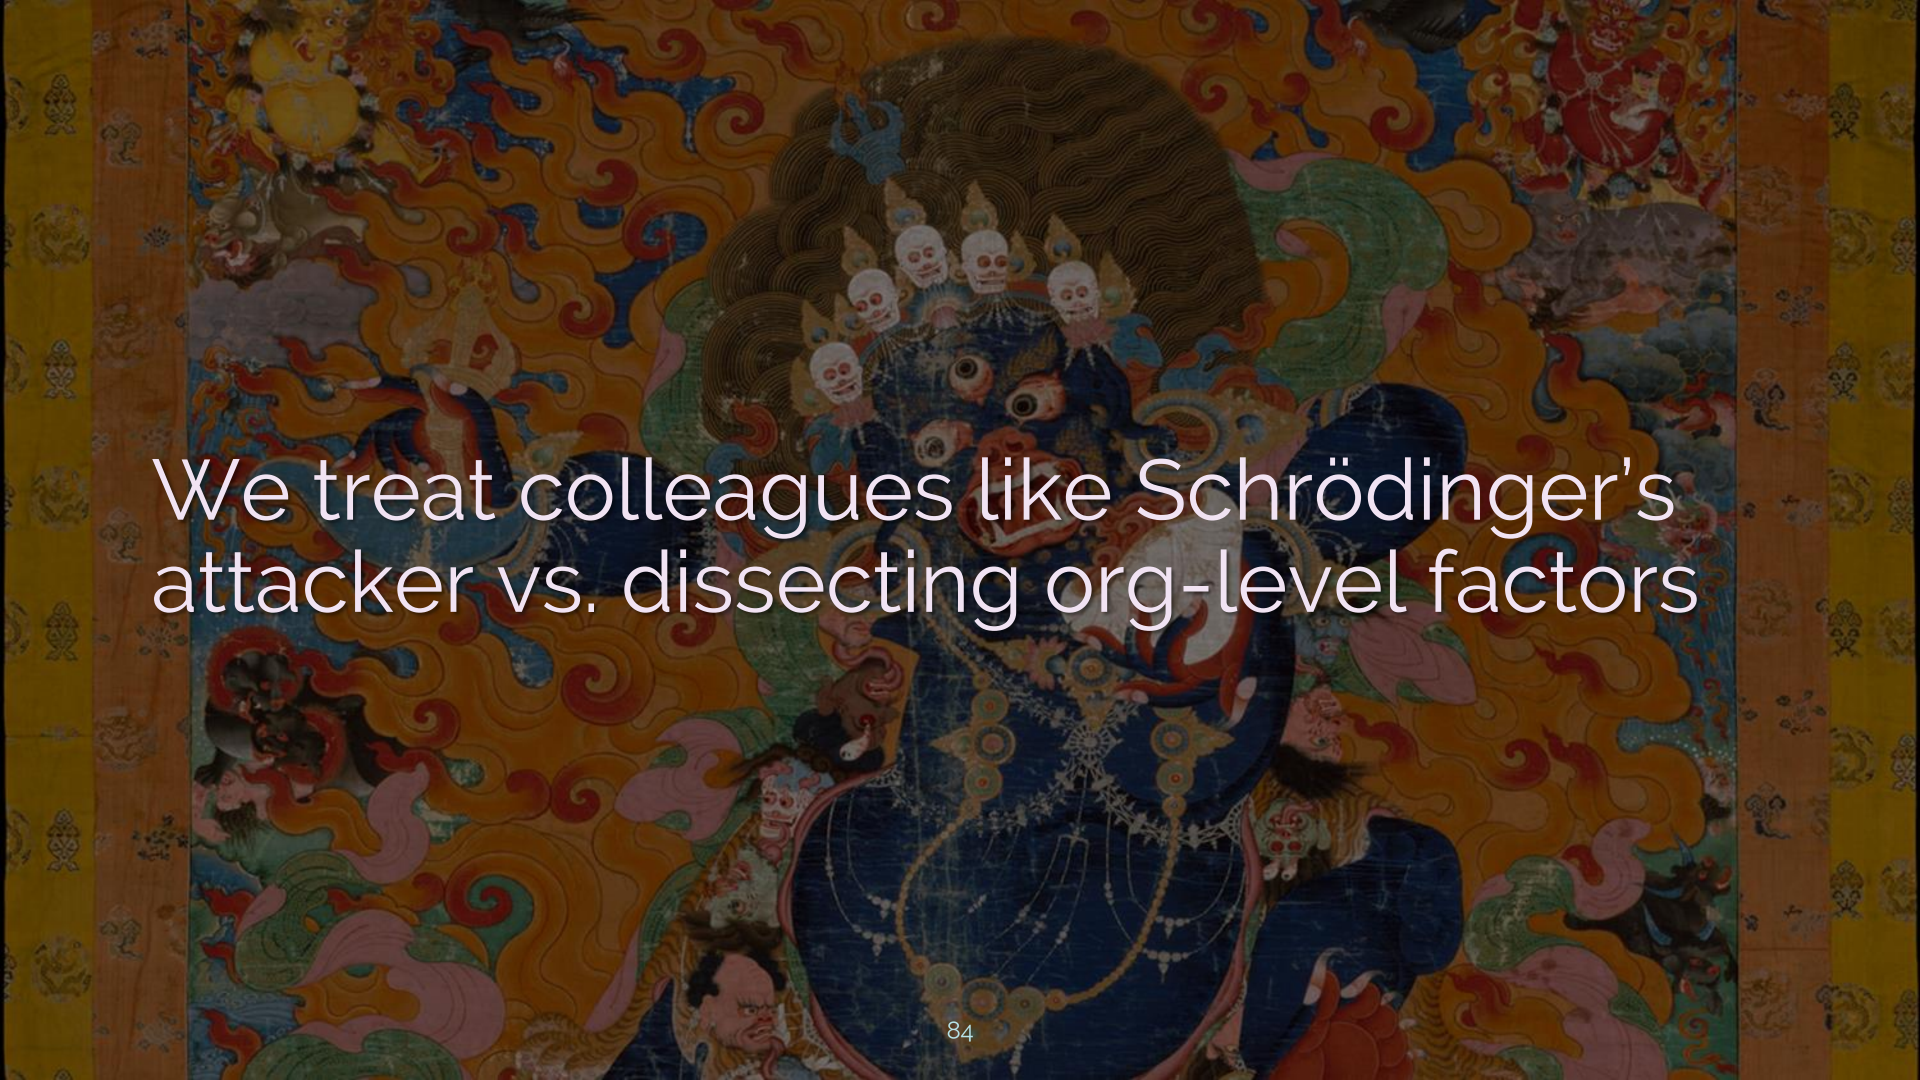

We treat colleagues like Schrödinger's attacker vs. dissecting org-level factors

Security is something a system does, not something a system has.

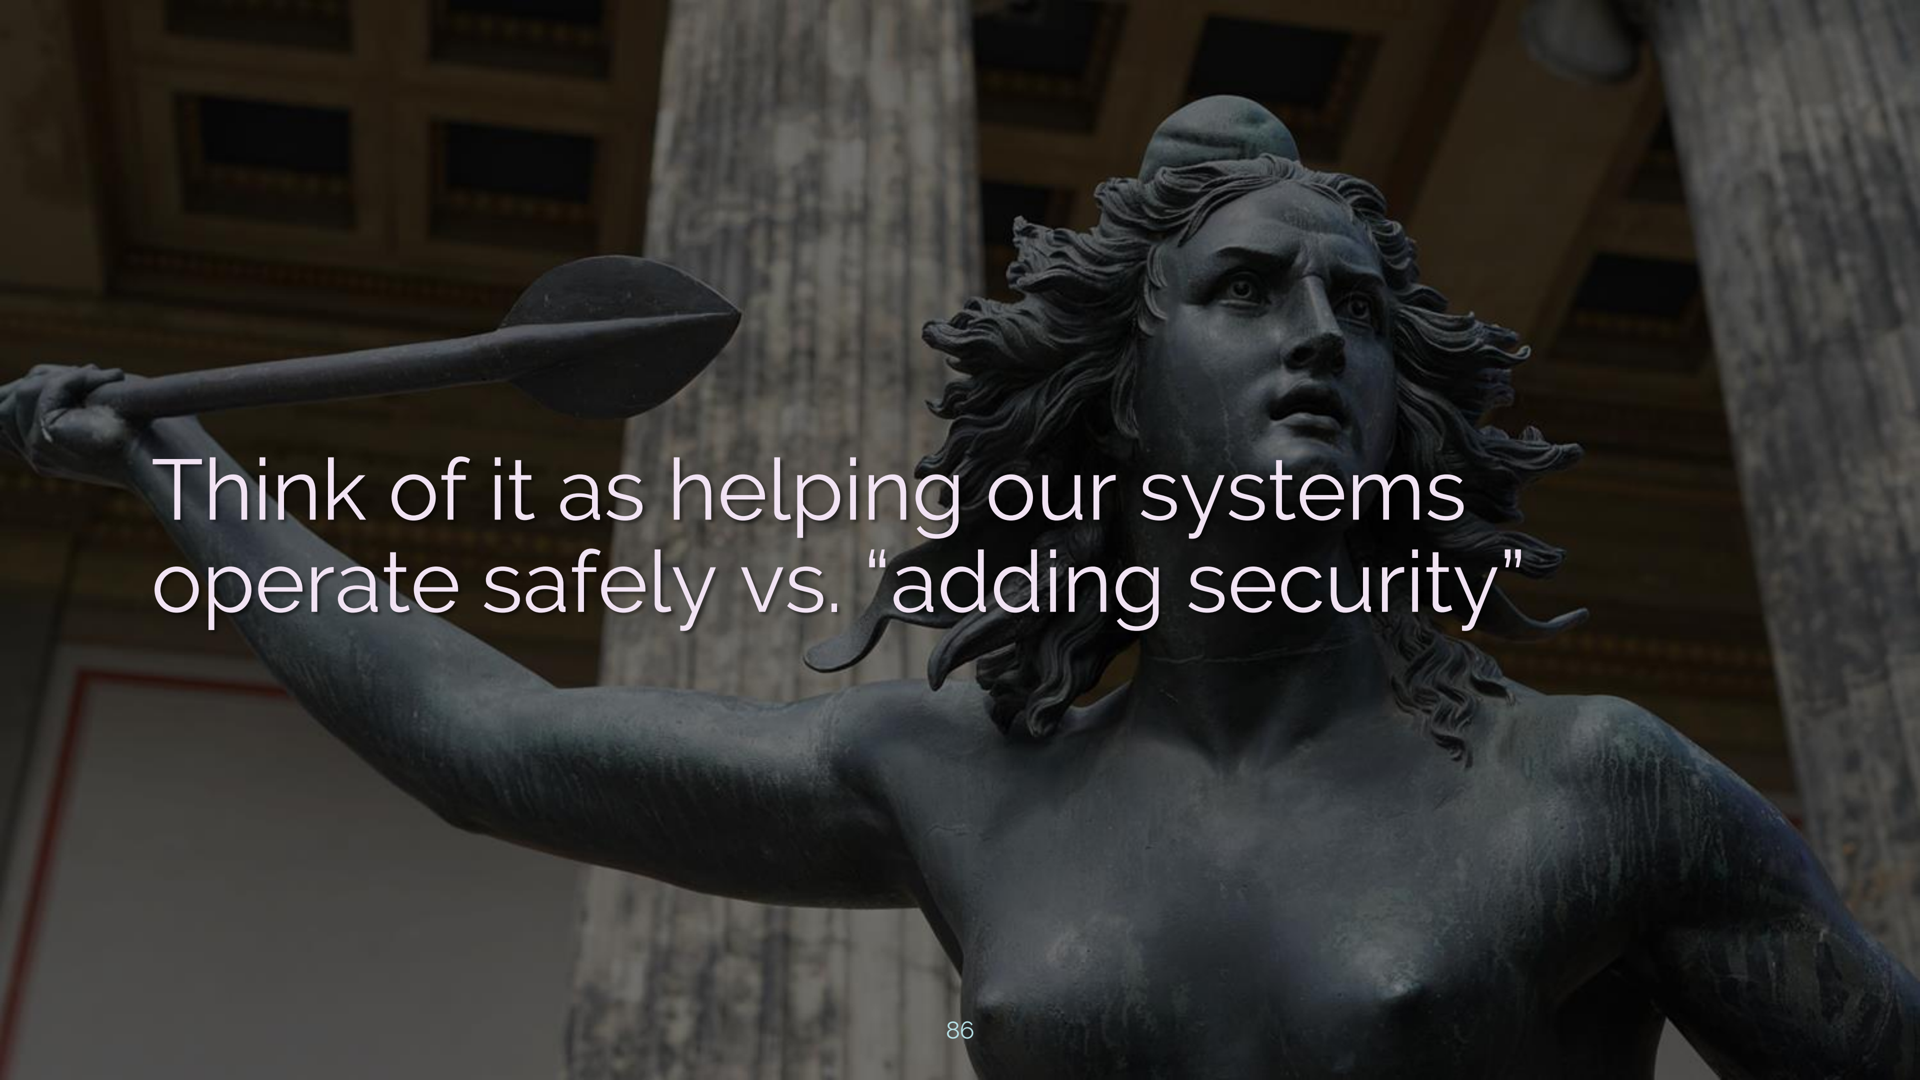

Think of it as helping our systems operate safely vs. “adding security”

Health & “security vanity” metrics don’t say whether systems are *doing* security

Number of vulns found matters less than their severity & how quickly they’re fixed

Infosec should analyze the mismatch between self-perception & reality

Alternative analysis for defenders is basically just user research…

Security UX

The pressure to meet competing goals is a strong source of security failure

What drives their promotion or firing?
What are their performance goals?

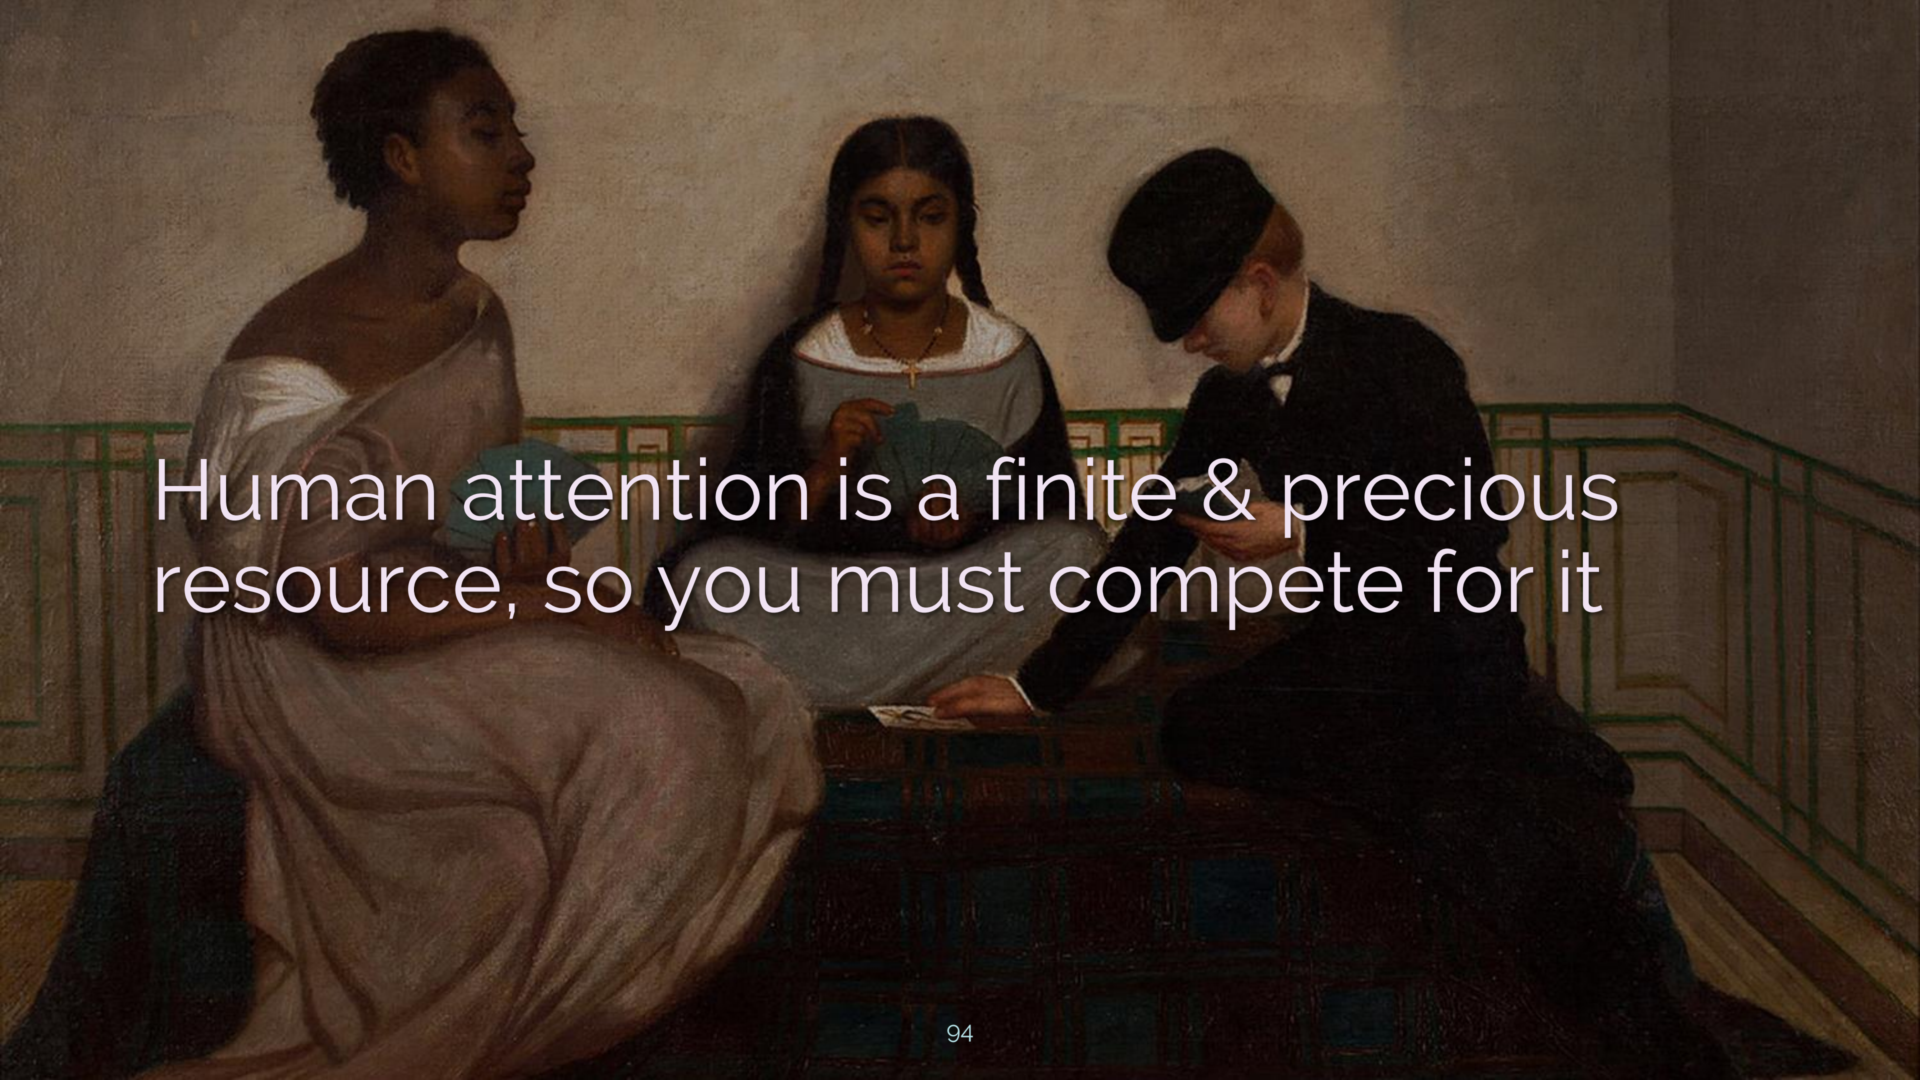

Human attention is a finite & precious resource, so you must compete for it

User research can help you determine how to draw attention towards security

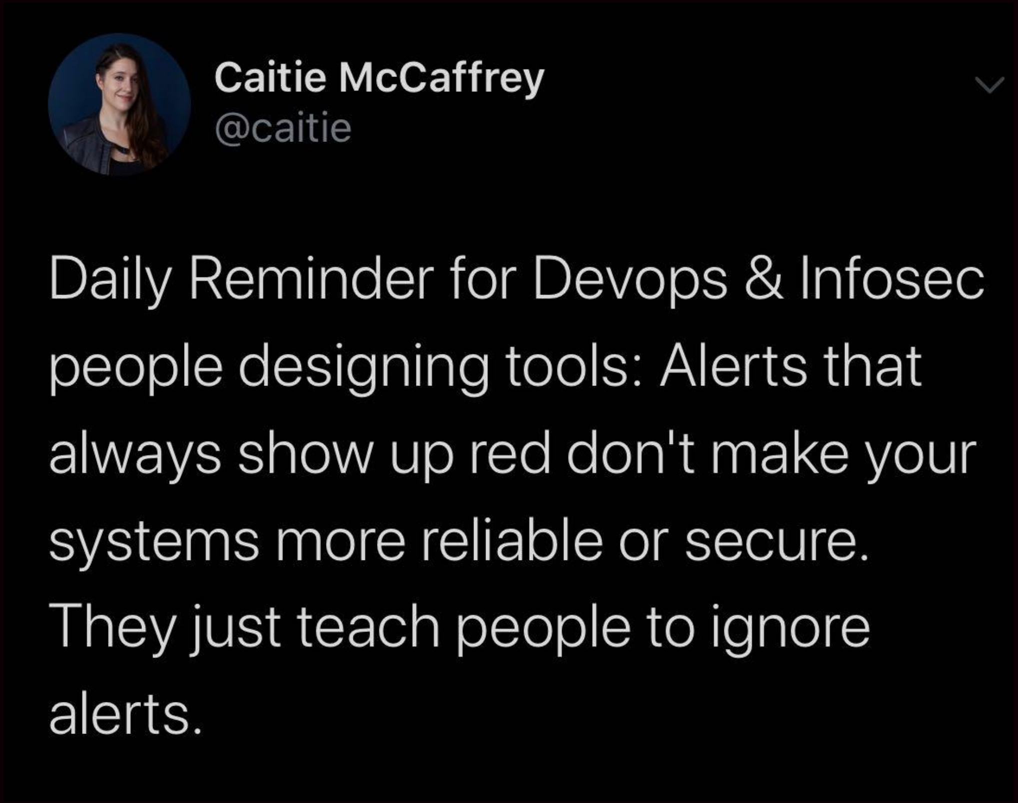

Caitie McCaffrey
@caitie

Daily Reminder for Devops & Infosec people designing tools: Alerts that always show up red don't make your systems more reliable or secure. They just teach people to ignore alerts.

WARNING: CYBER ANOMALY

(thanks Raytheon)

Choice architecture: organizing the context in which people make decisions

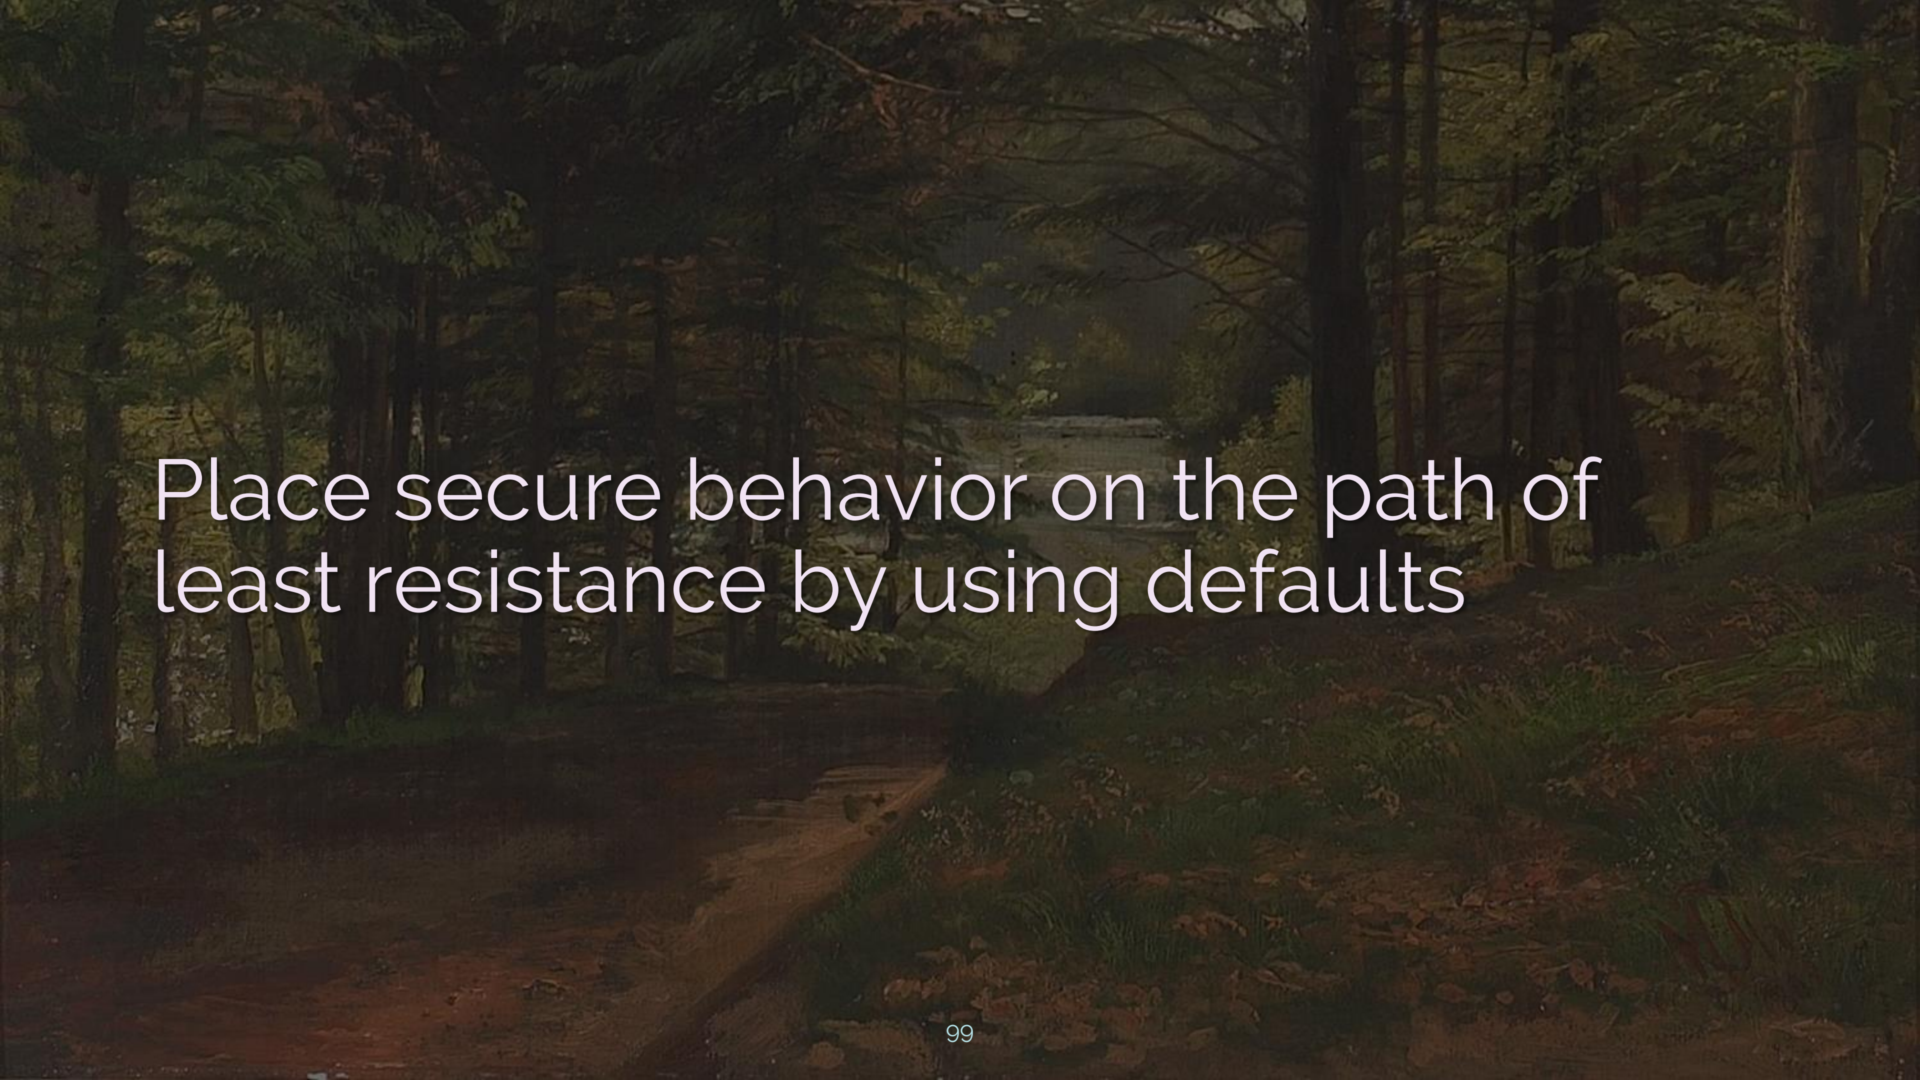

Place secure behavior on the path of least resistance by using defaults

e.g. Requiring 2FA to create an account, security tests in CI/CD pipelines

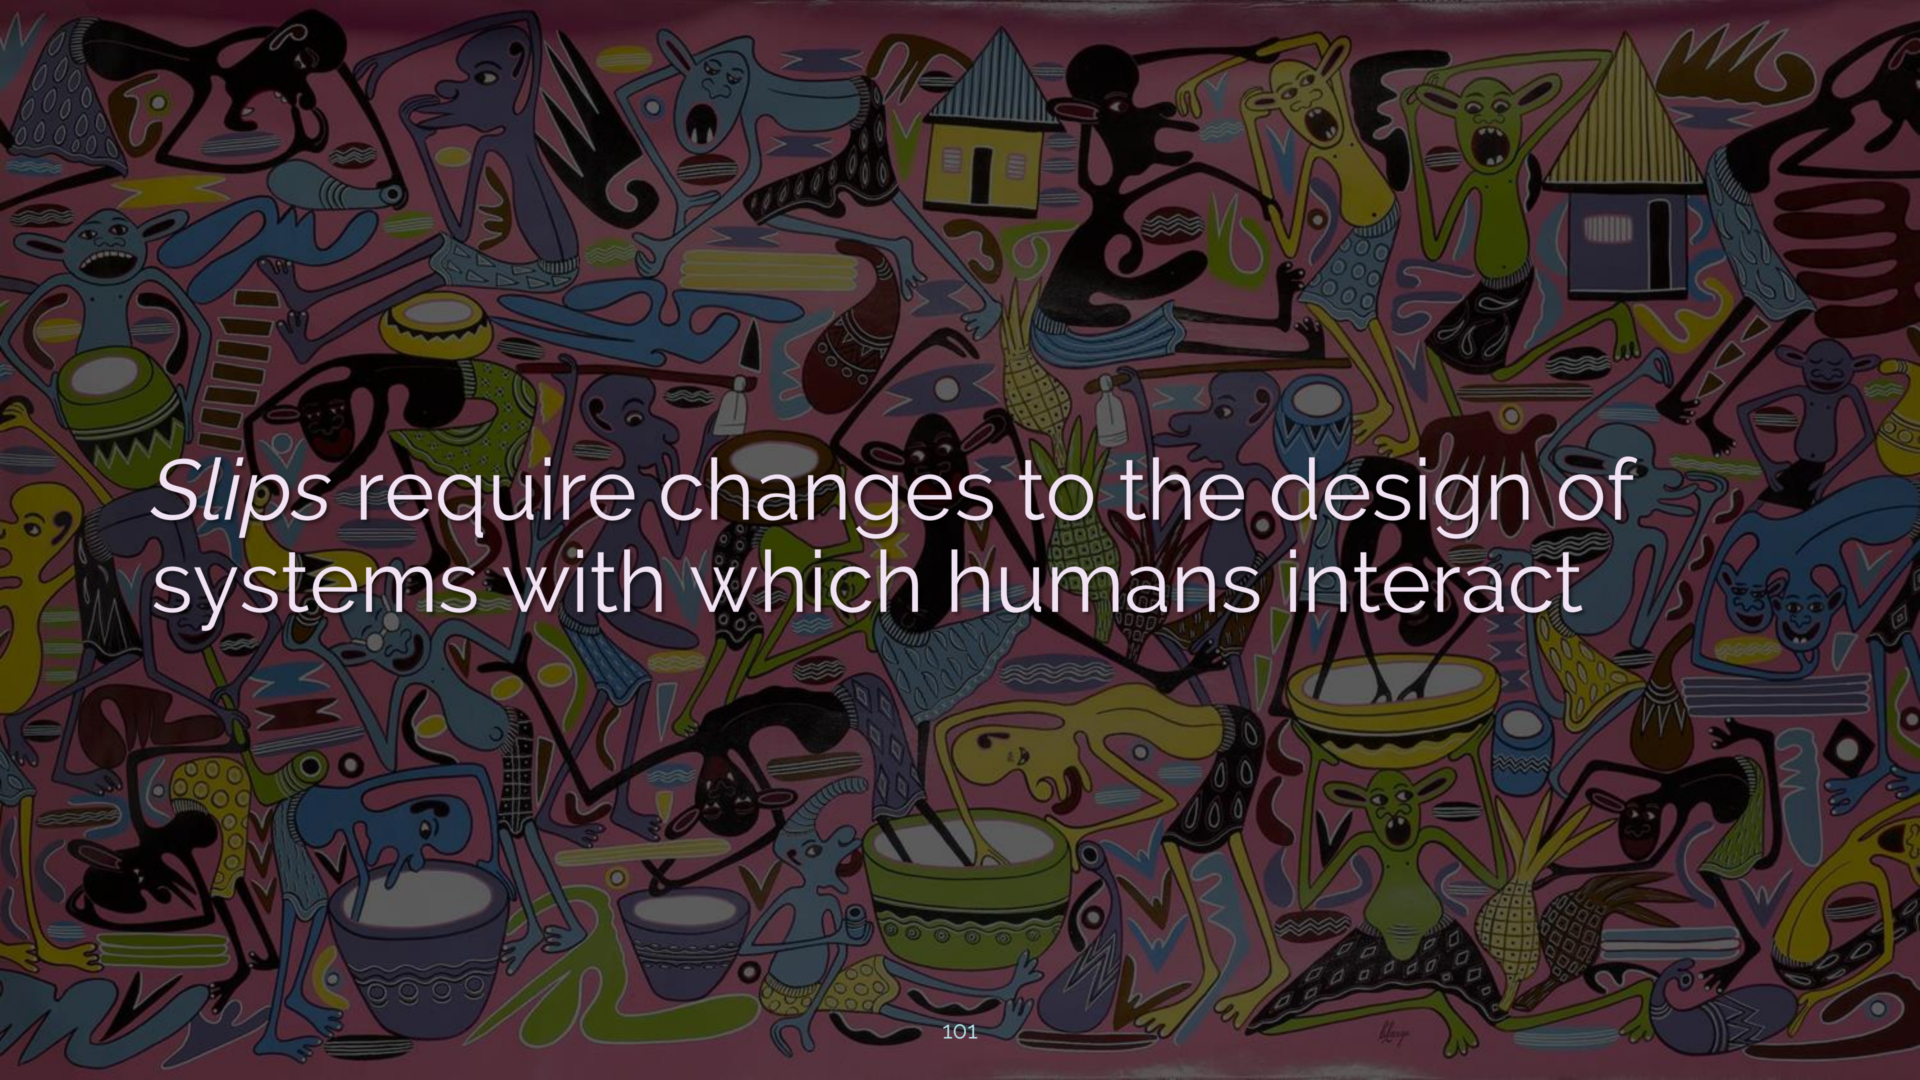

Slips require changes to the design of systems with which humans interact

Checklists, defaults, eliminating distractions, removing complexity…

Strong security design anticipates user workarounds & safely supports them

e.g. Self-service app approvals with a Slackbot to confirm the run request

Think in terms of acceptable tradeoffs – create secure alternatives, not loopholes

How else can you better understand your systems & the context they create?

Chaos Security Engineering

We will never be able to eliminate the potential for error.

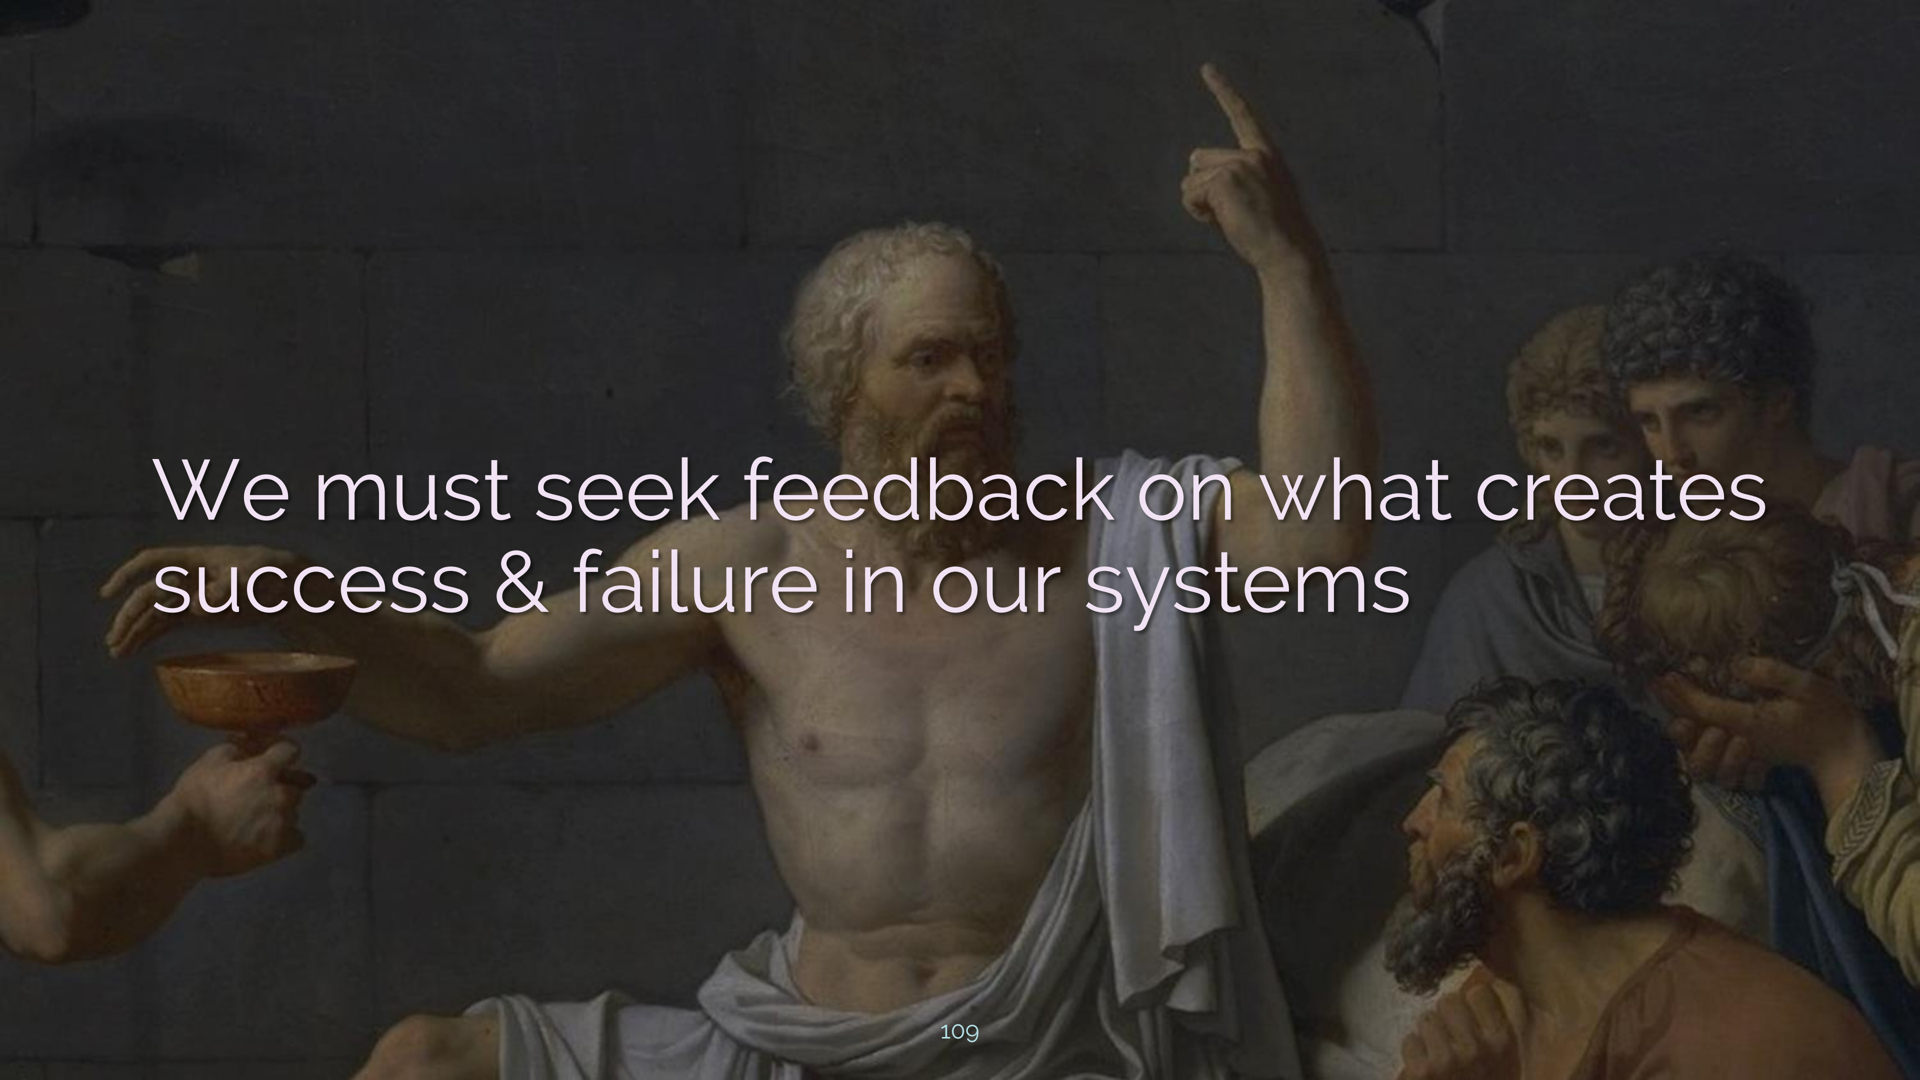

We must seek feedback on what creates success & failure in our systems

“Enhancing error tolerance, error detection, and error recovery together produce safety.”

– Woods, et al

Error tolerance: the ability to not get totally pwned when compromise occurs

Error detection: the ability to spot unwanted activity

Error recovery: the ability to restore systems to their intended functionality

Highest ROI: anticipating how the potential for failure evolves

Chaos eng: continual experimentation to evaluate response to unexpected failure

e.g. Retrograding: inject old versions of libs, containers, etc. into your systems

Chaos engineering assumes existing knowledge hangs in a delicate balance

The potential for hazard is constantly changing, creating new blindspots

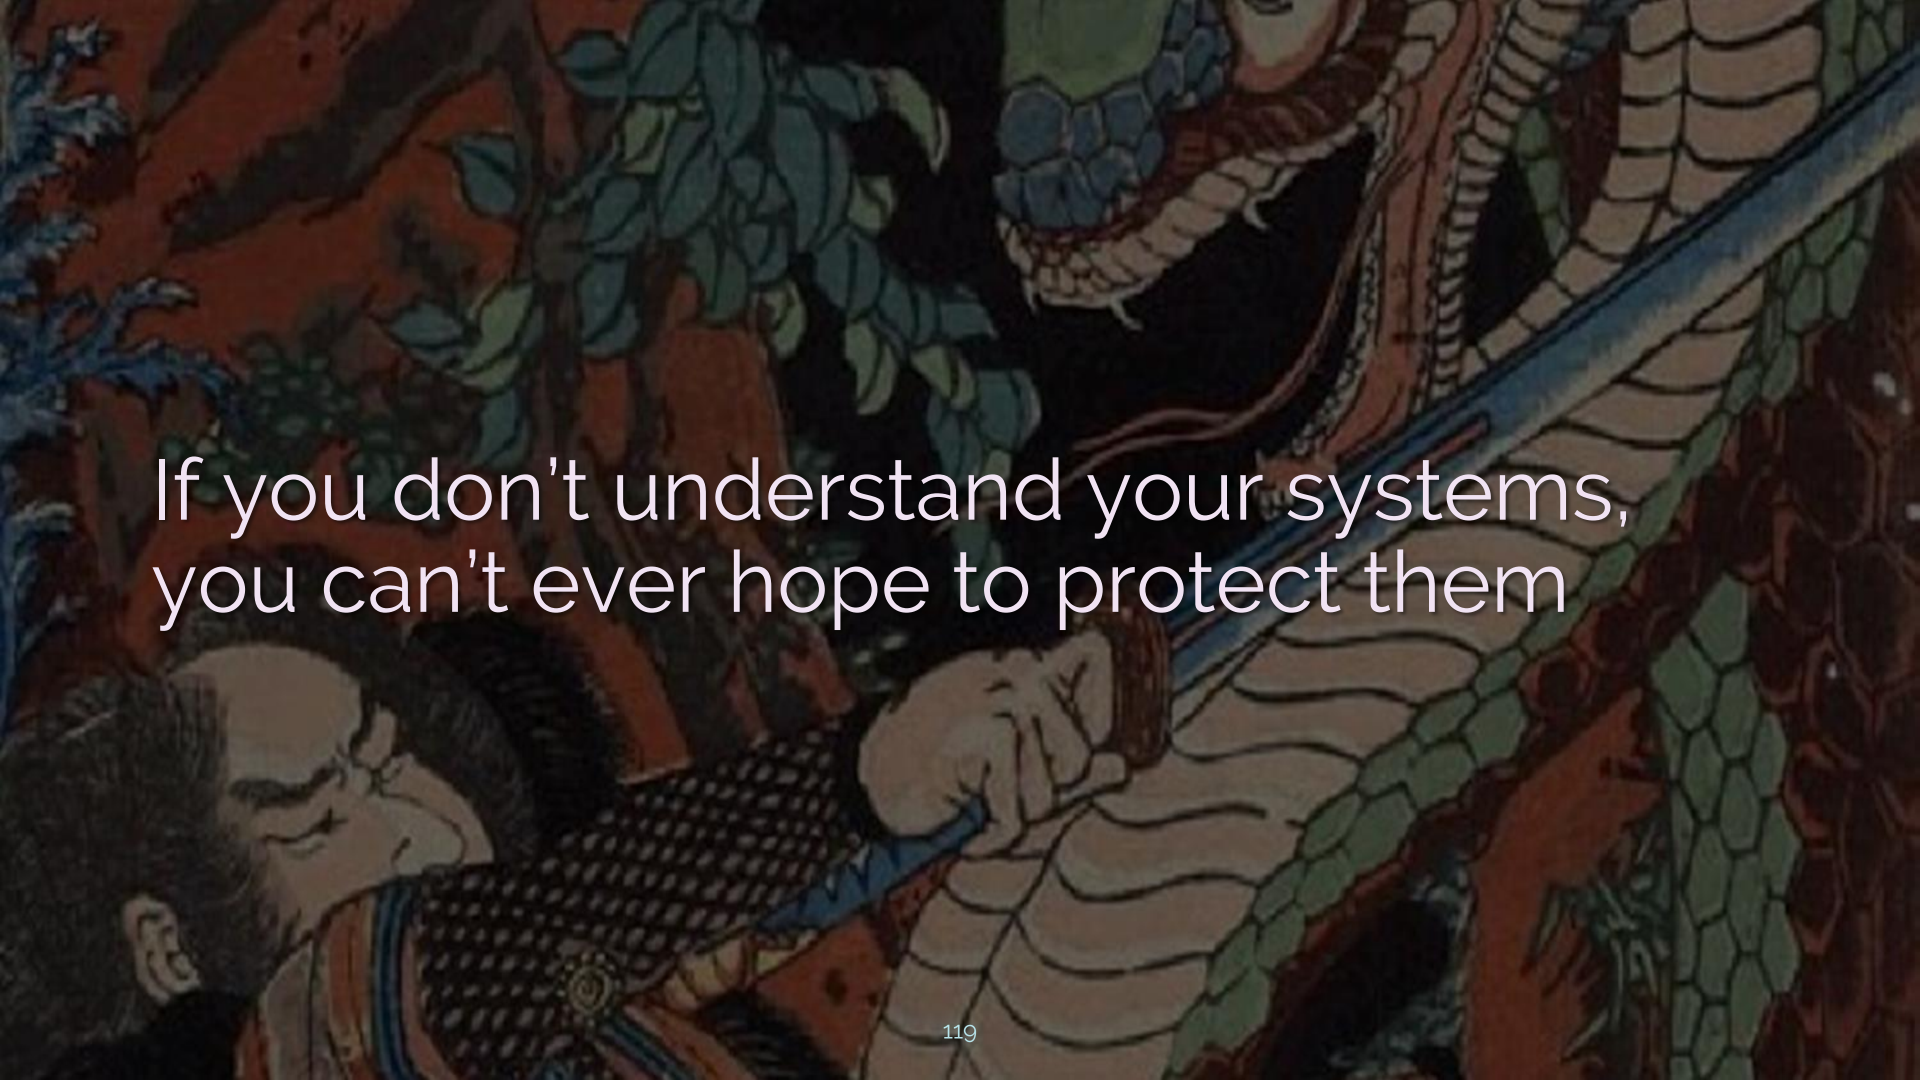

If you don’t understand your systems, you can’t ever hope to protect them

Chaos security engineering requires a blameless culture…

Blameless Culture

A blameless culture balances safety and accountability – not absolution

Supports a perpetual state of learning, in which critical info isn't suppressed

Asking the right questions is the first step towards a blameless culture

Neutral questions prevent bias from seeping into our incident review

Ask other practitioners what they would do in the same original context

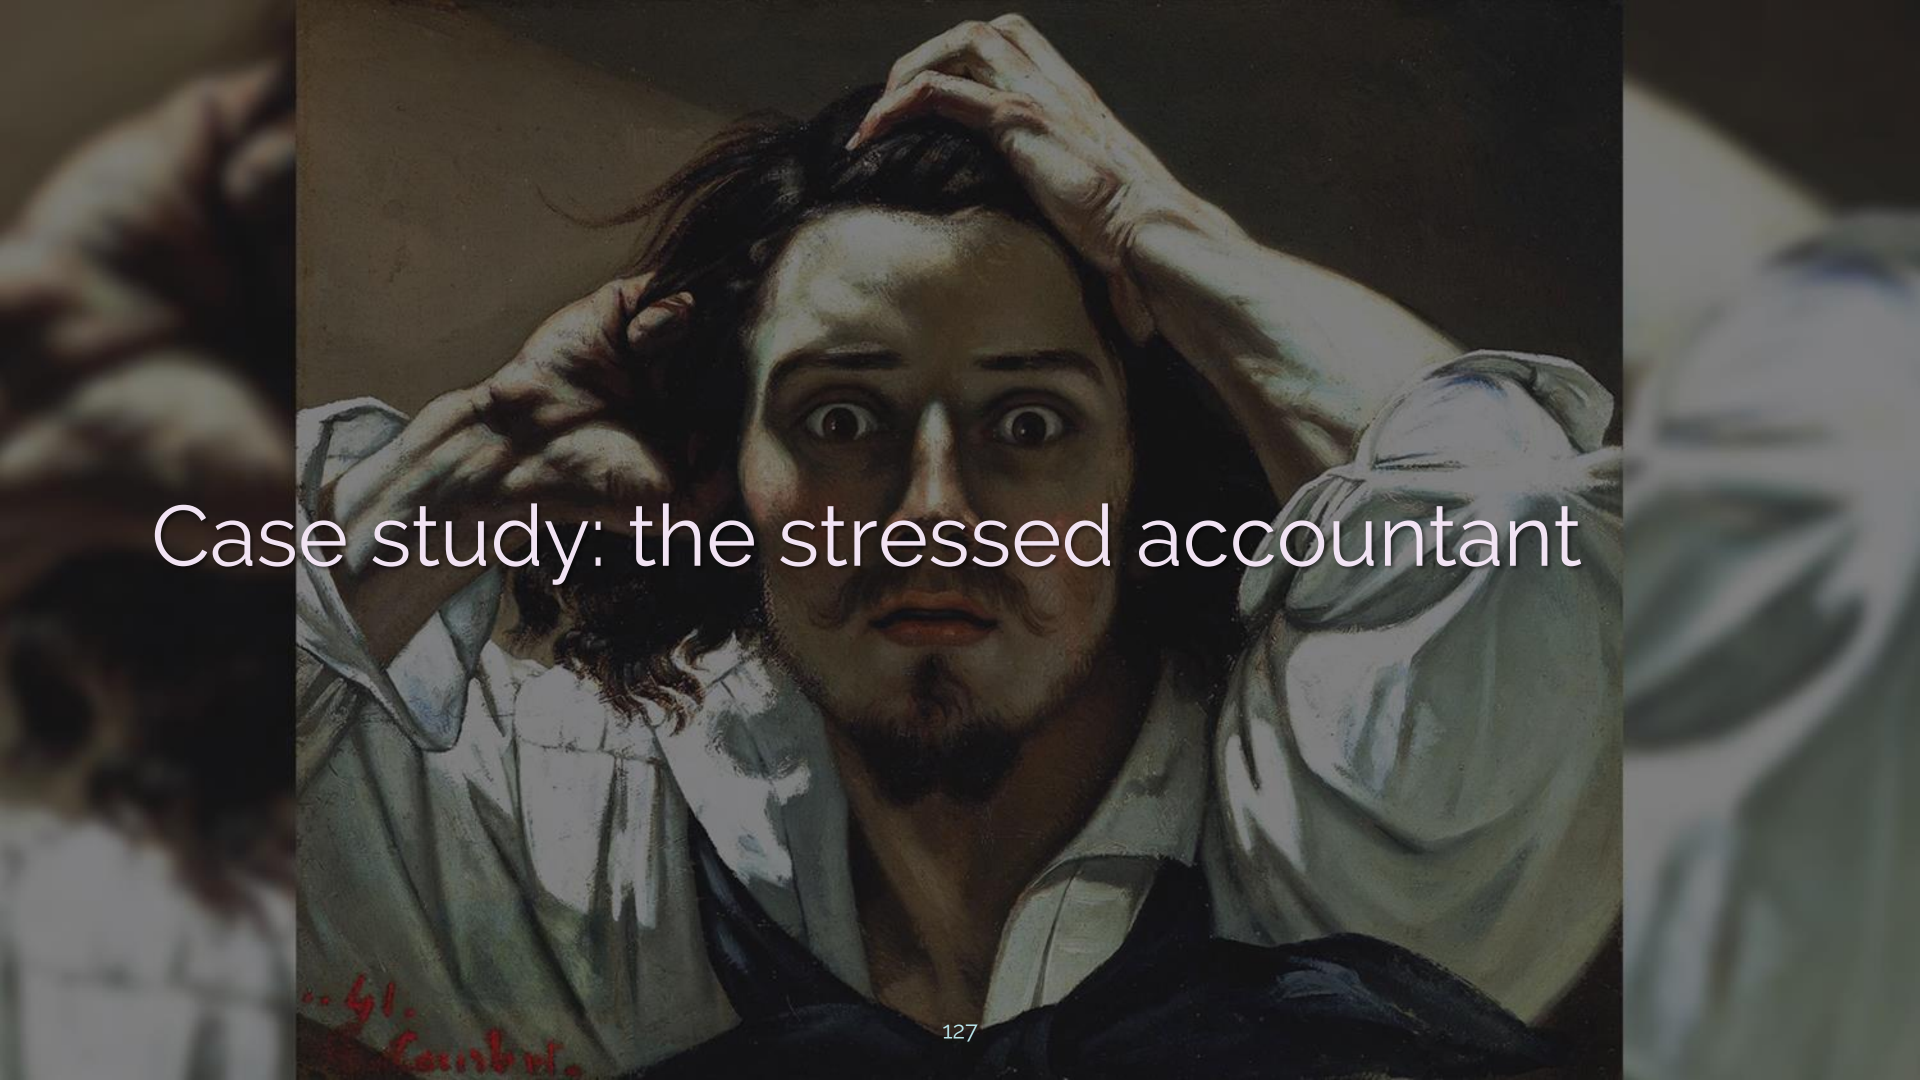

Case study: the stressed accountant

“Human error” becomes a reasonable action given the human’s circumstances

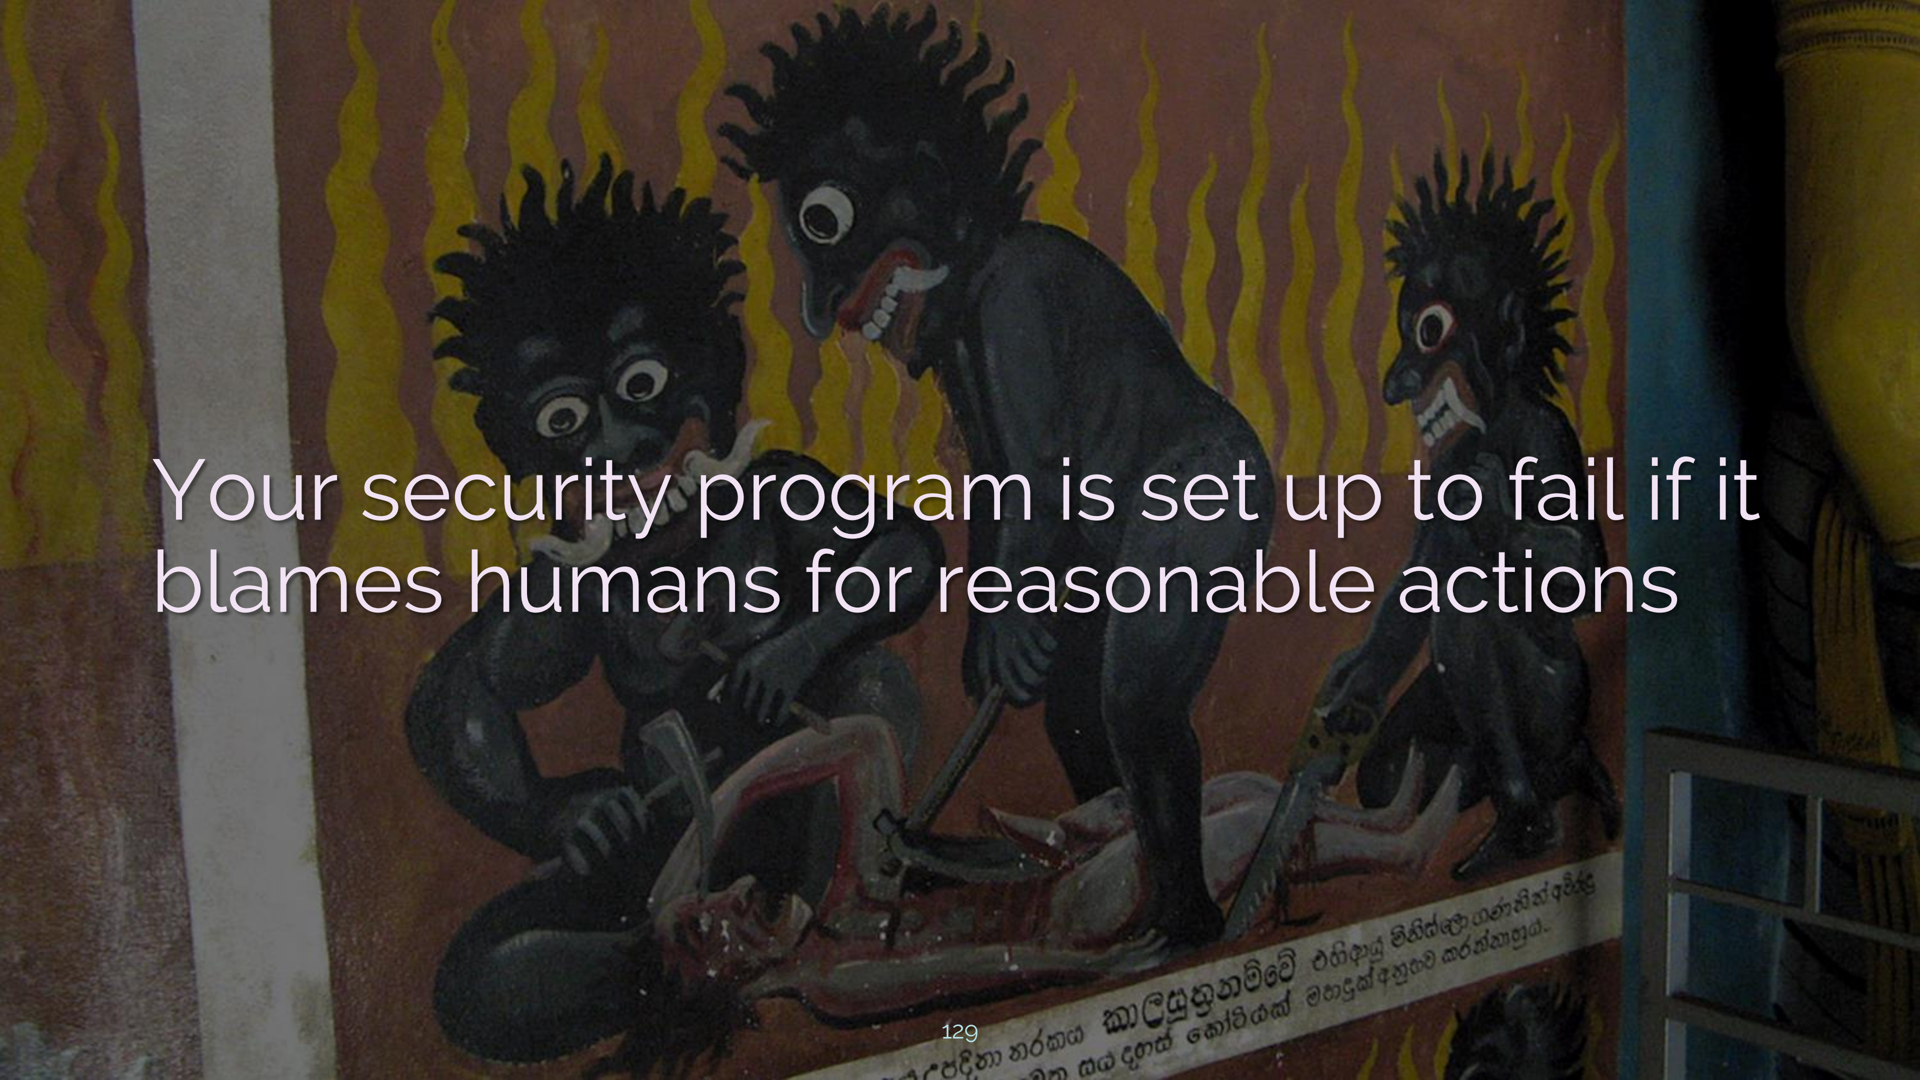
උපදිනා නරකය කාලසූත්‍රනම්වේ එහිආයු මිනිස්ලොව ගණනින් අවුරුදු
සය දහස් කෝටියක් මහාදුක් අනුභව කරන්නාහුය.

Your security program is set up to fail if it blames humans for reasonable actions

Neutral practitioner questions help sketch a portrait of local rationality

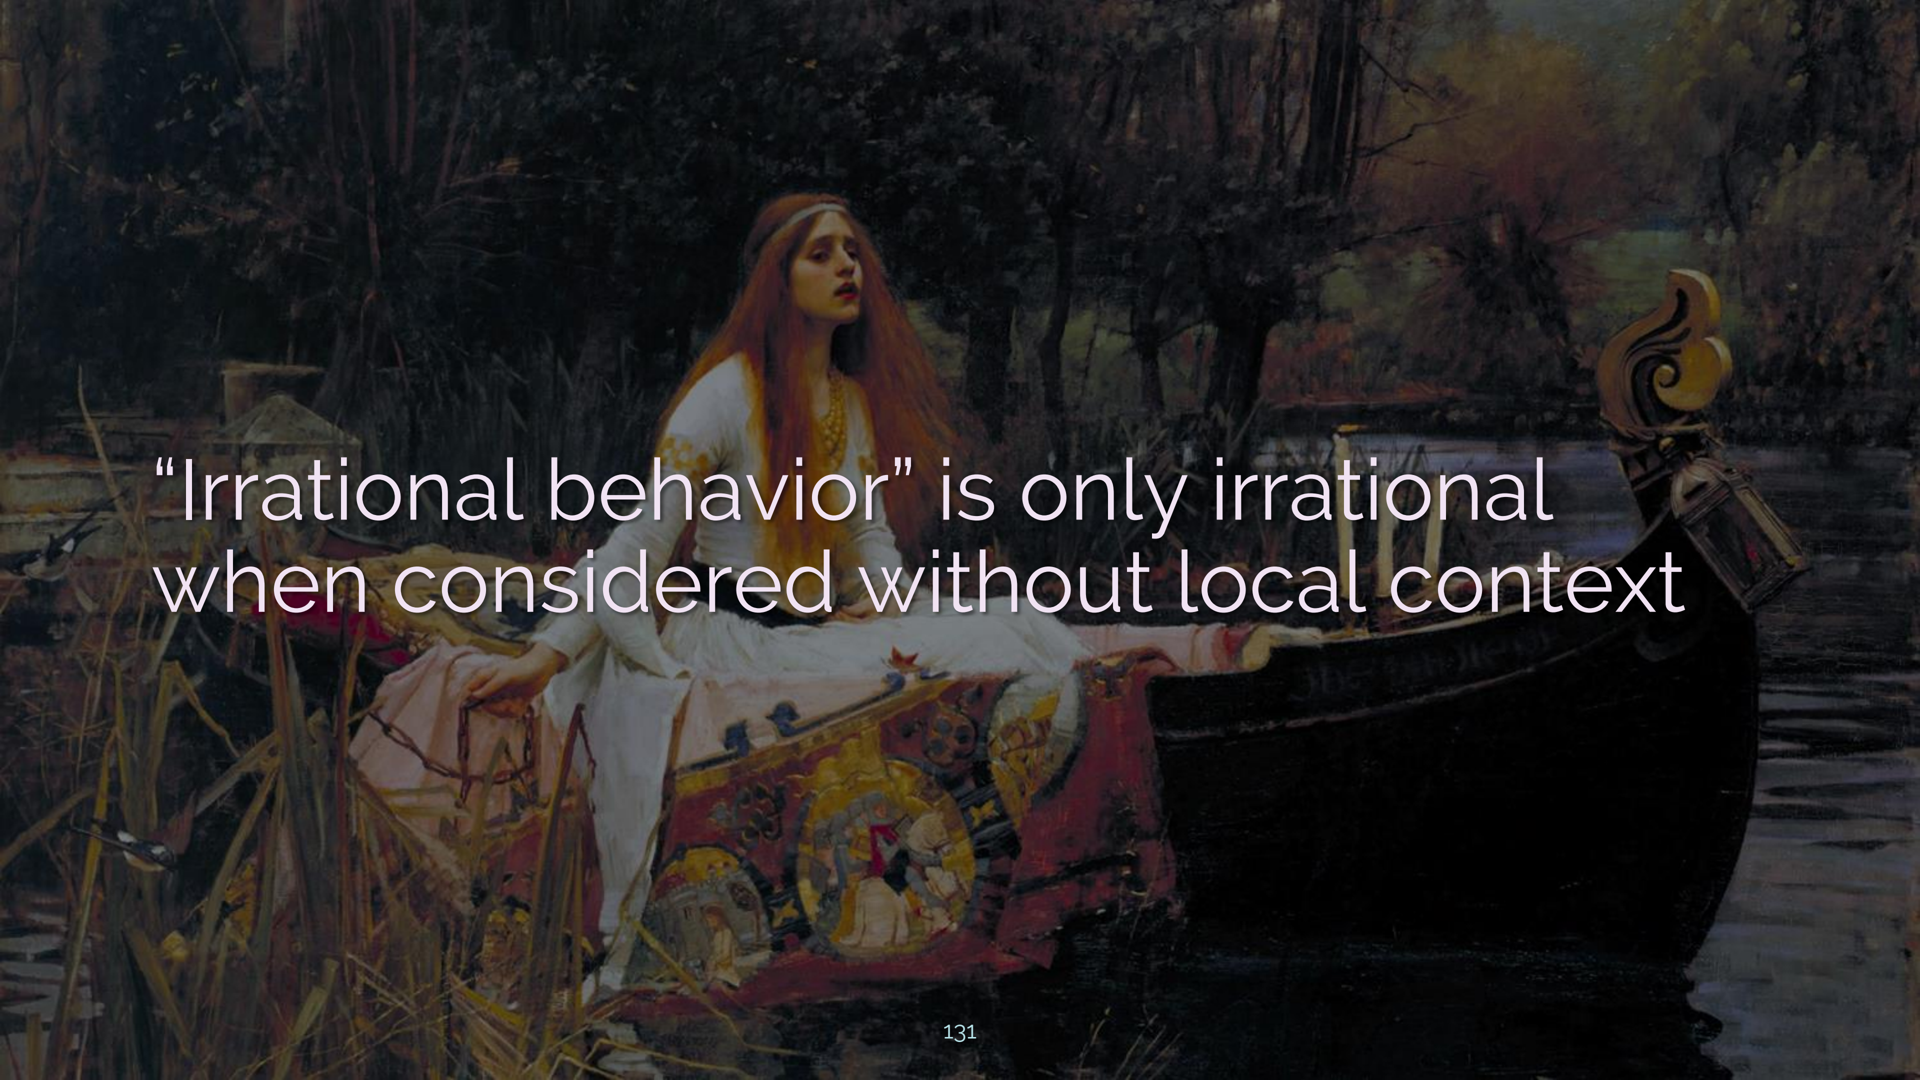

“Irrational behavior” is only irrational
when considered without local context

Our goal is to change the *context* of decision-making to promote security

If you’re using an ad-hominem attack in incident review, you’ve veered astray

In Conclusion

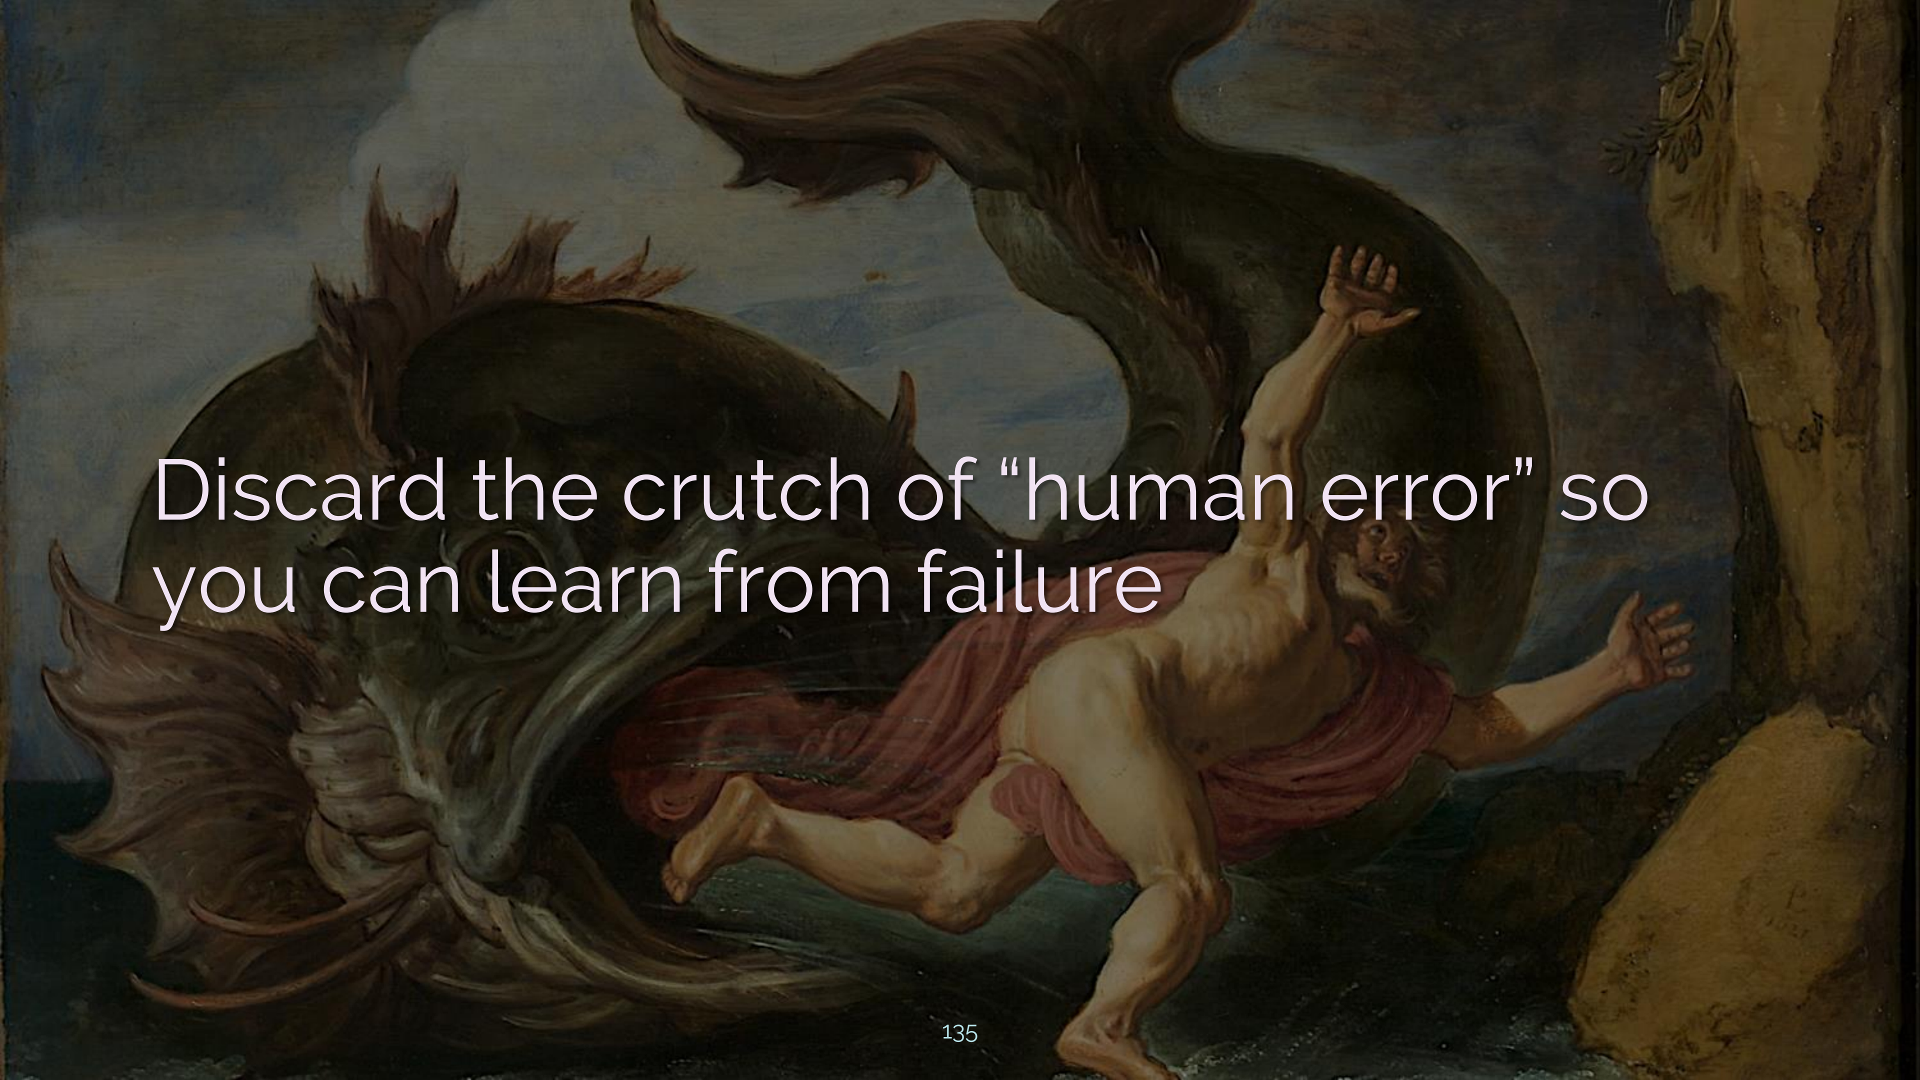

Discard the crutch of “human error” so you can learn from failure

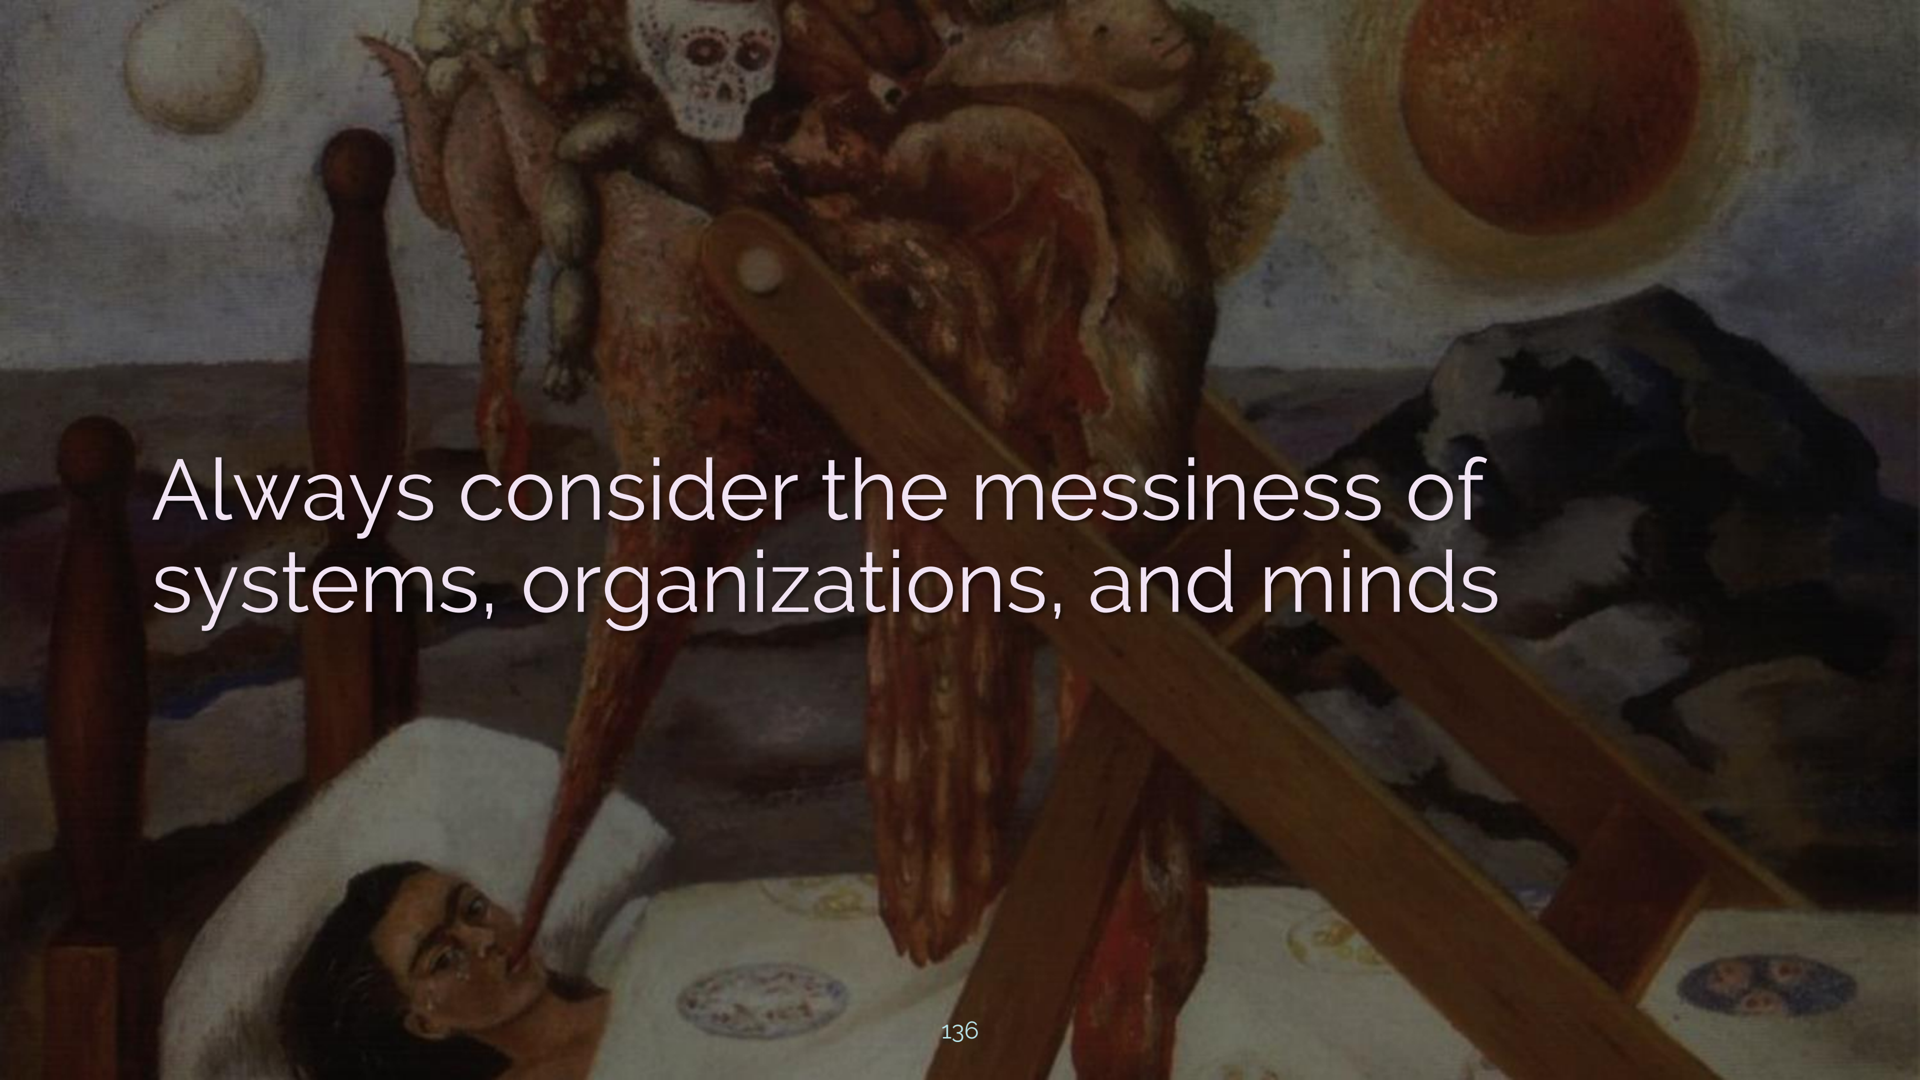

Always consider the messiness of systems, organizations, and minds

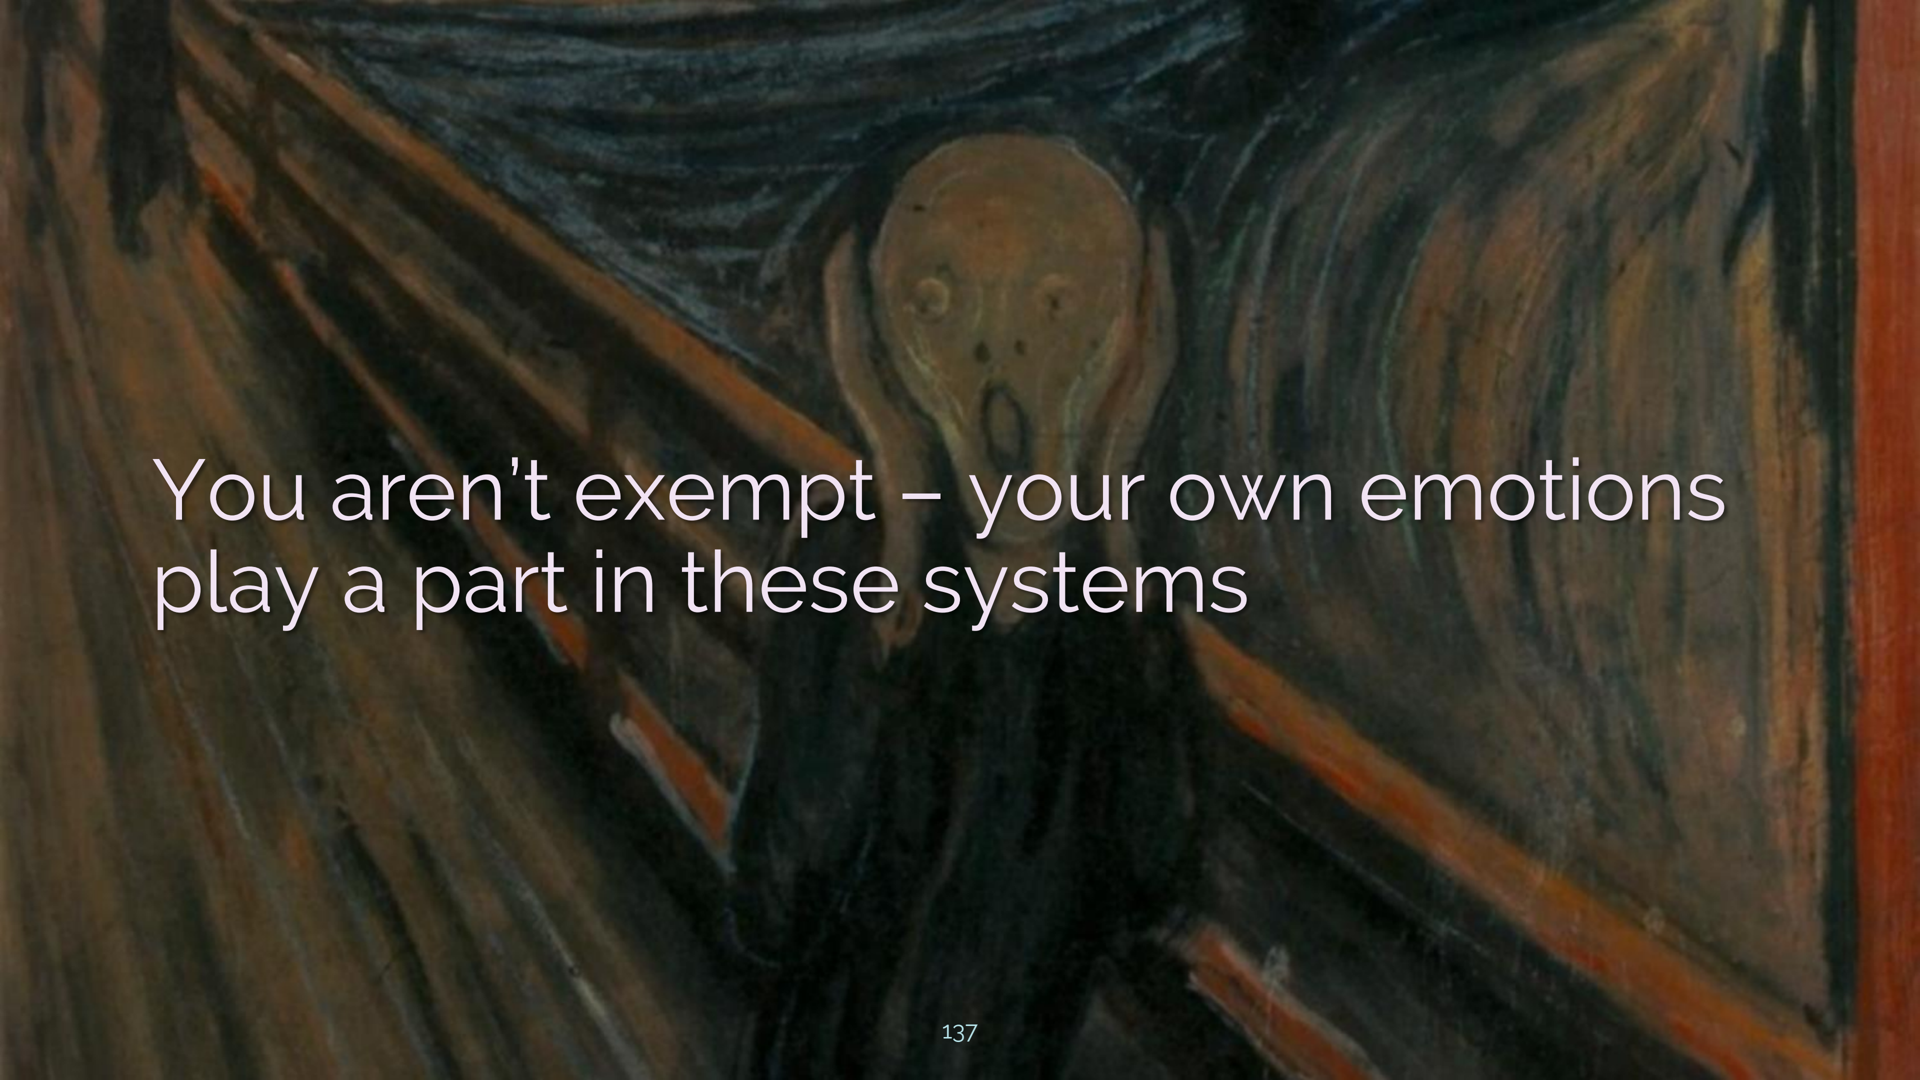

You aren’t exempt – your own emotions play a part in these systems

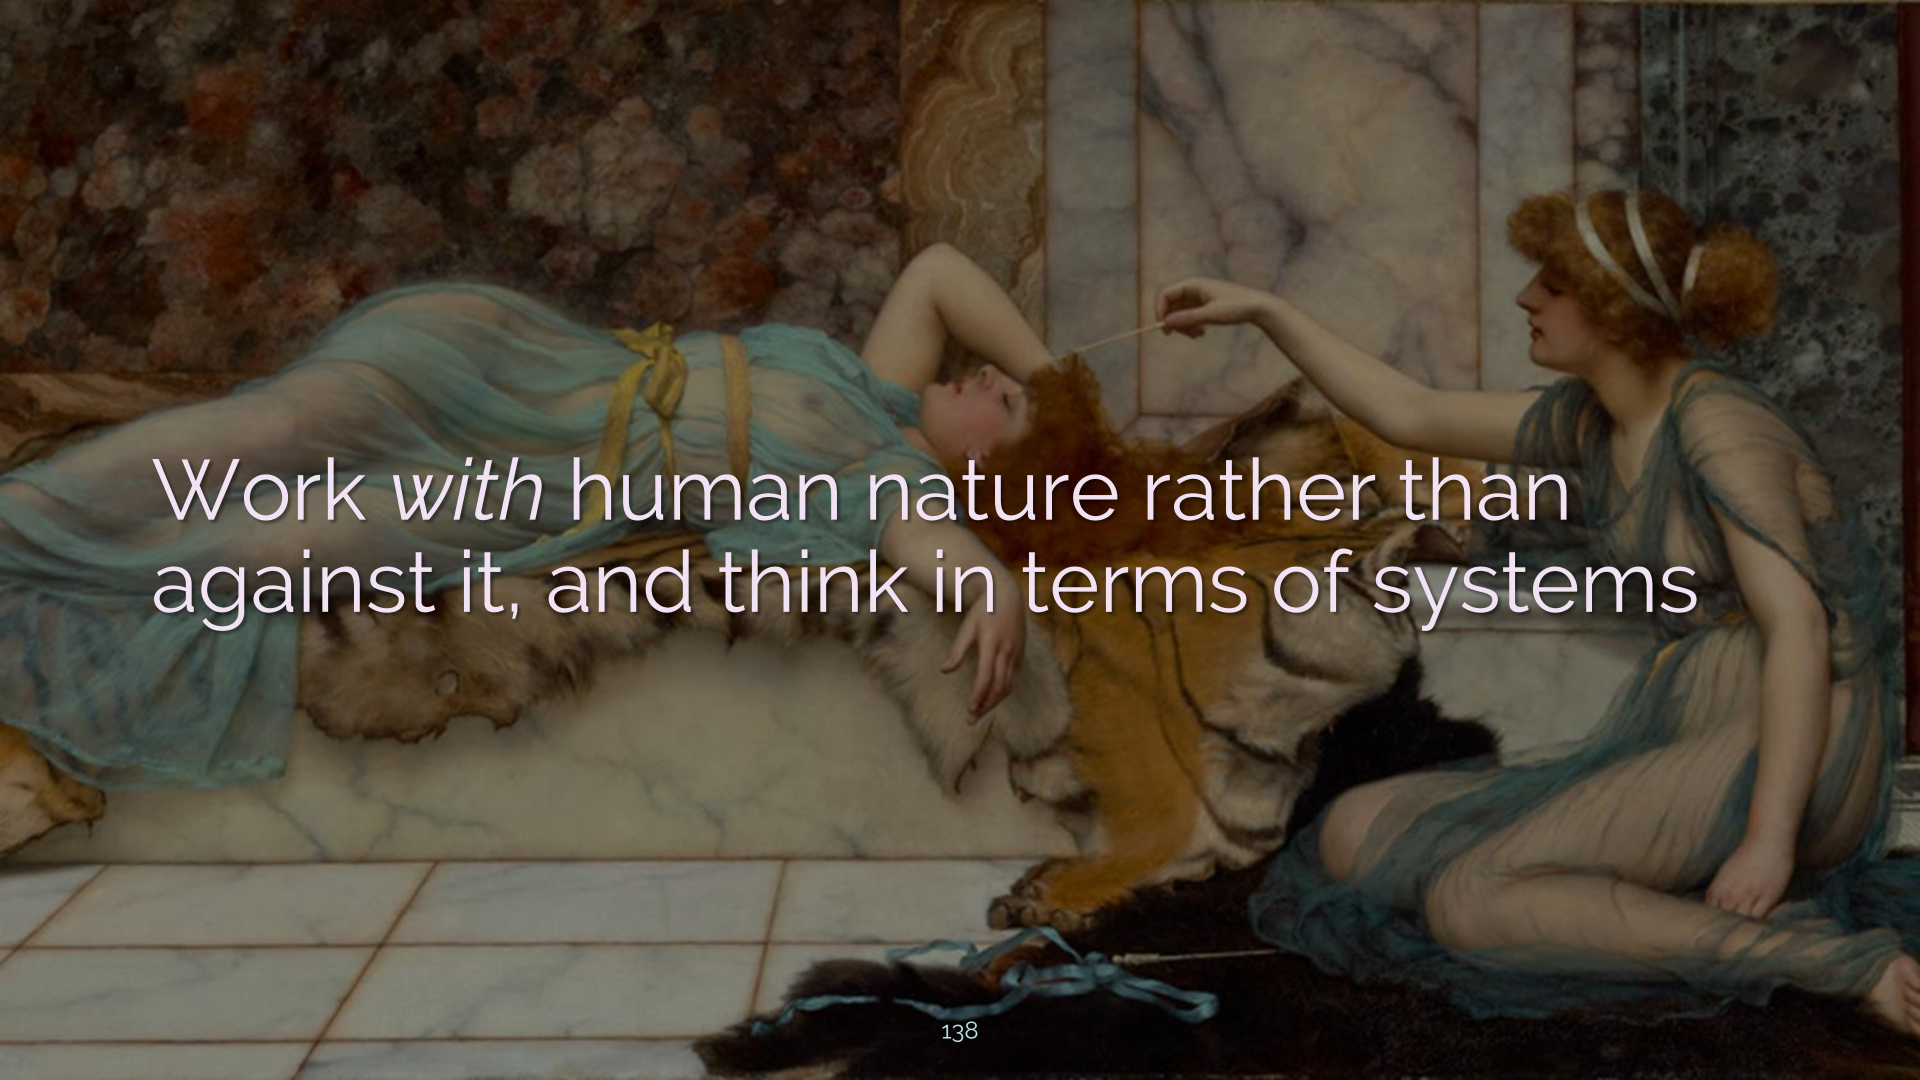

Work *with* human nature rather than against it, and think in terms of systems

Leverage UX & chaos eng to improve the context your systems engender

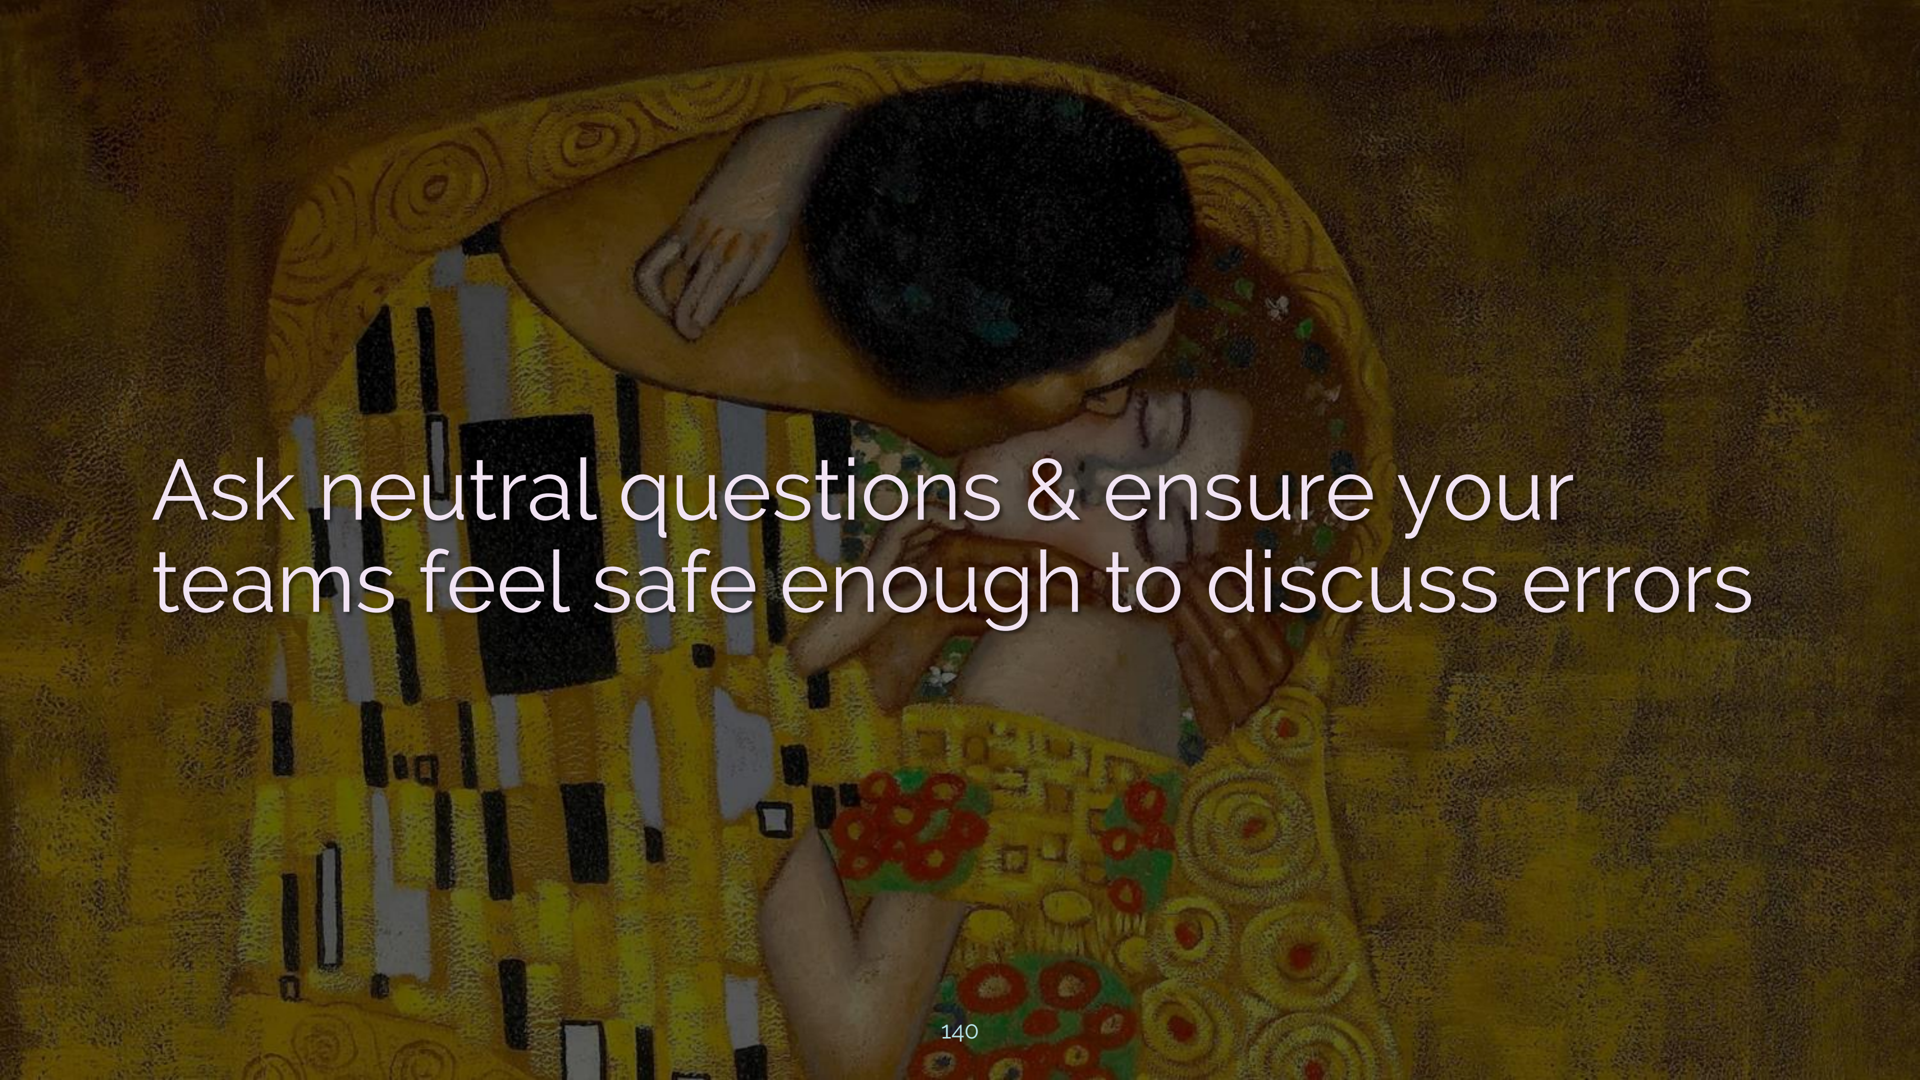

Ask neutral questions & ensure your teams feel safe enough to discuss errors

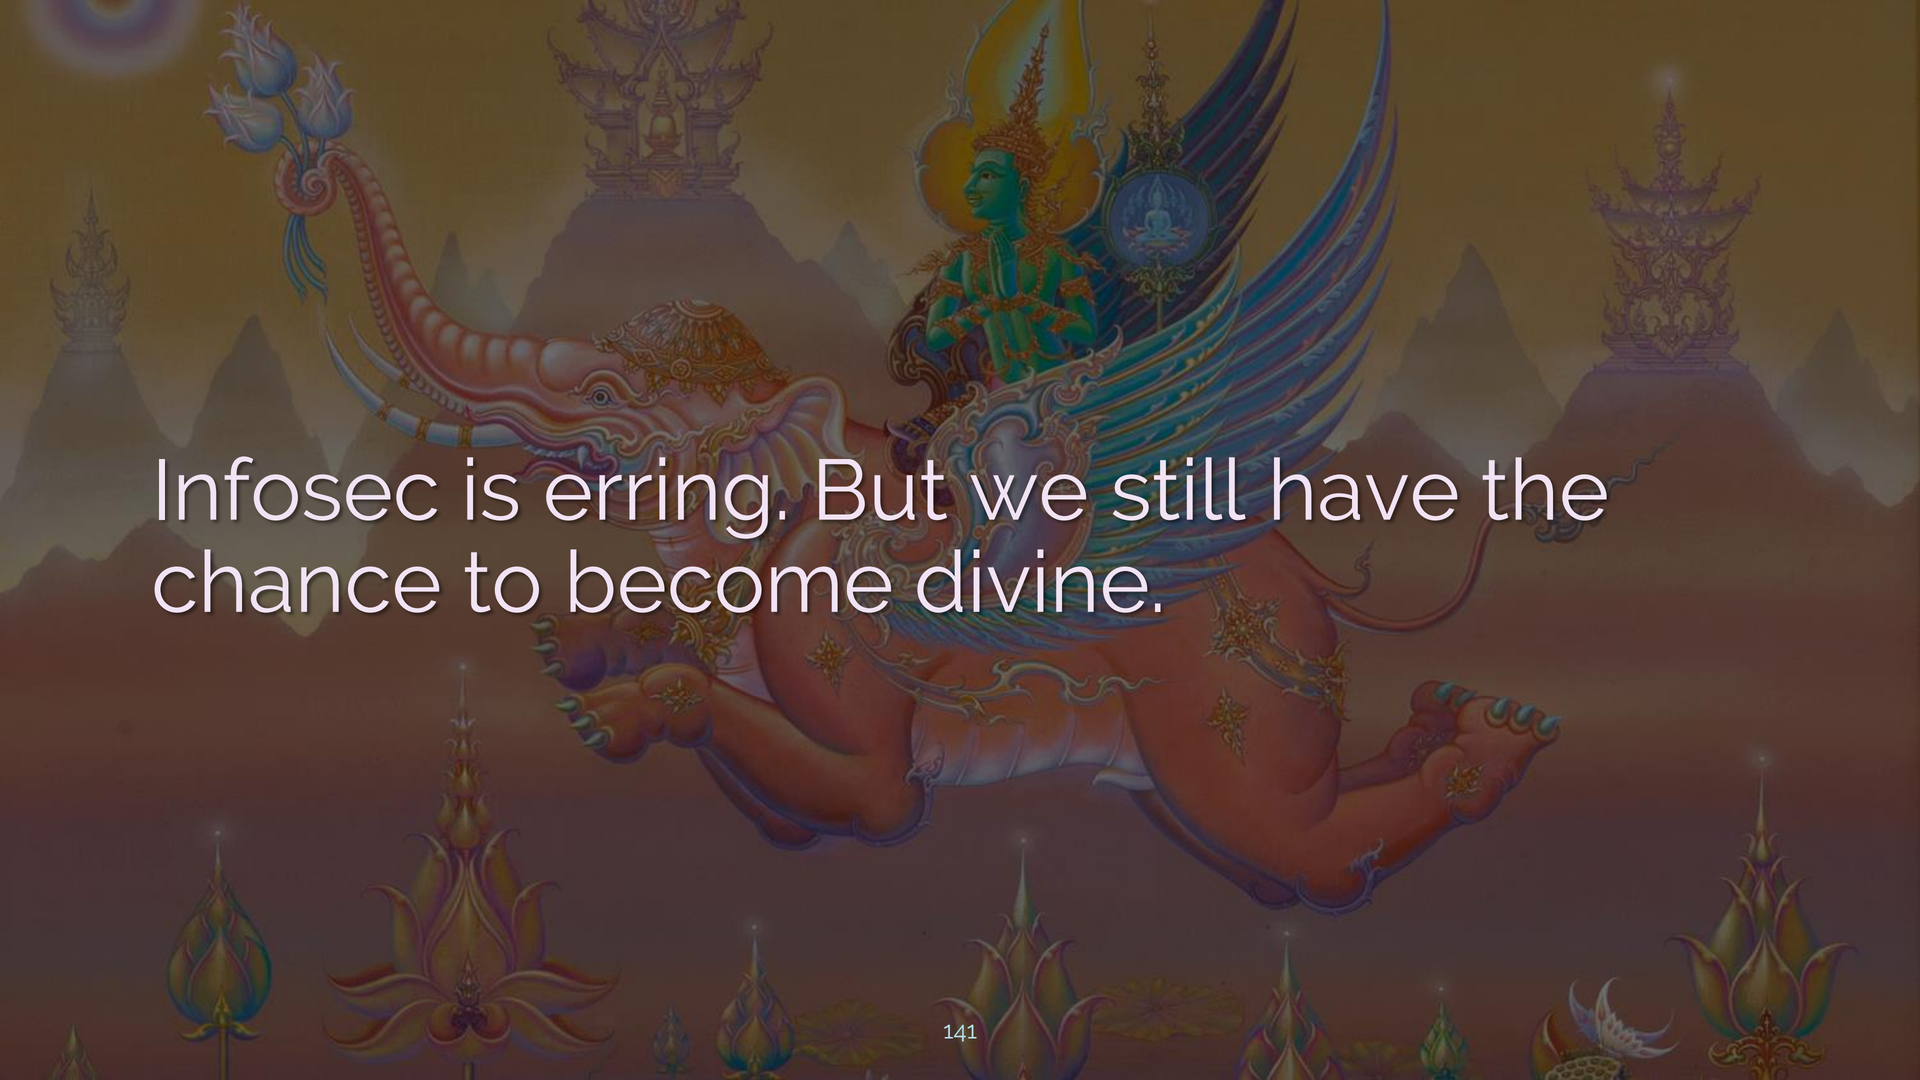

Infosec is erring. But we still have the chance to become divine.

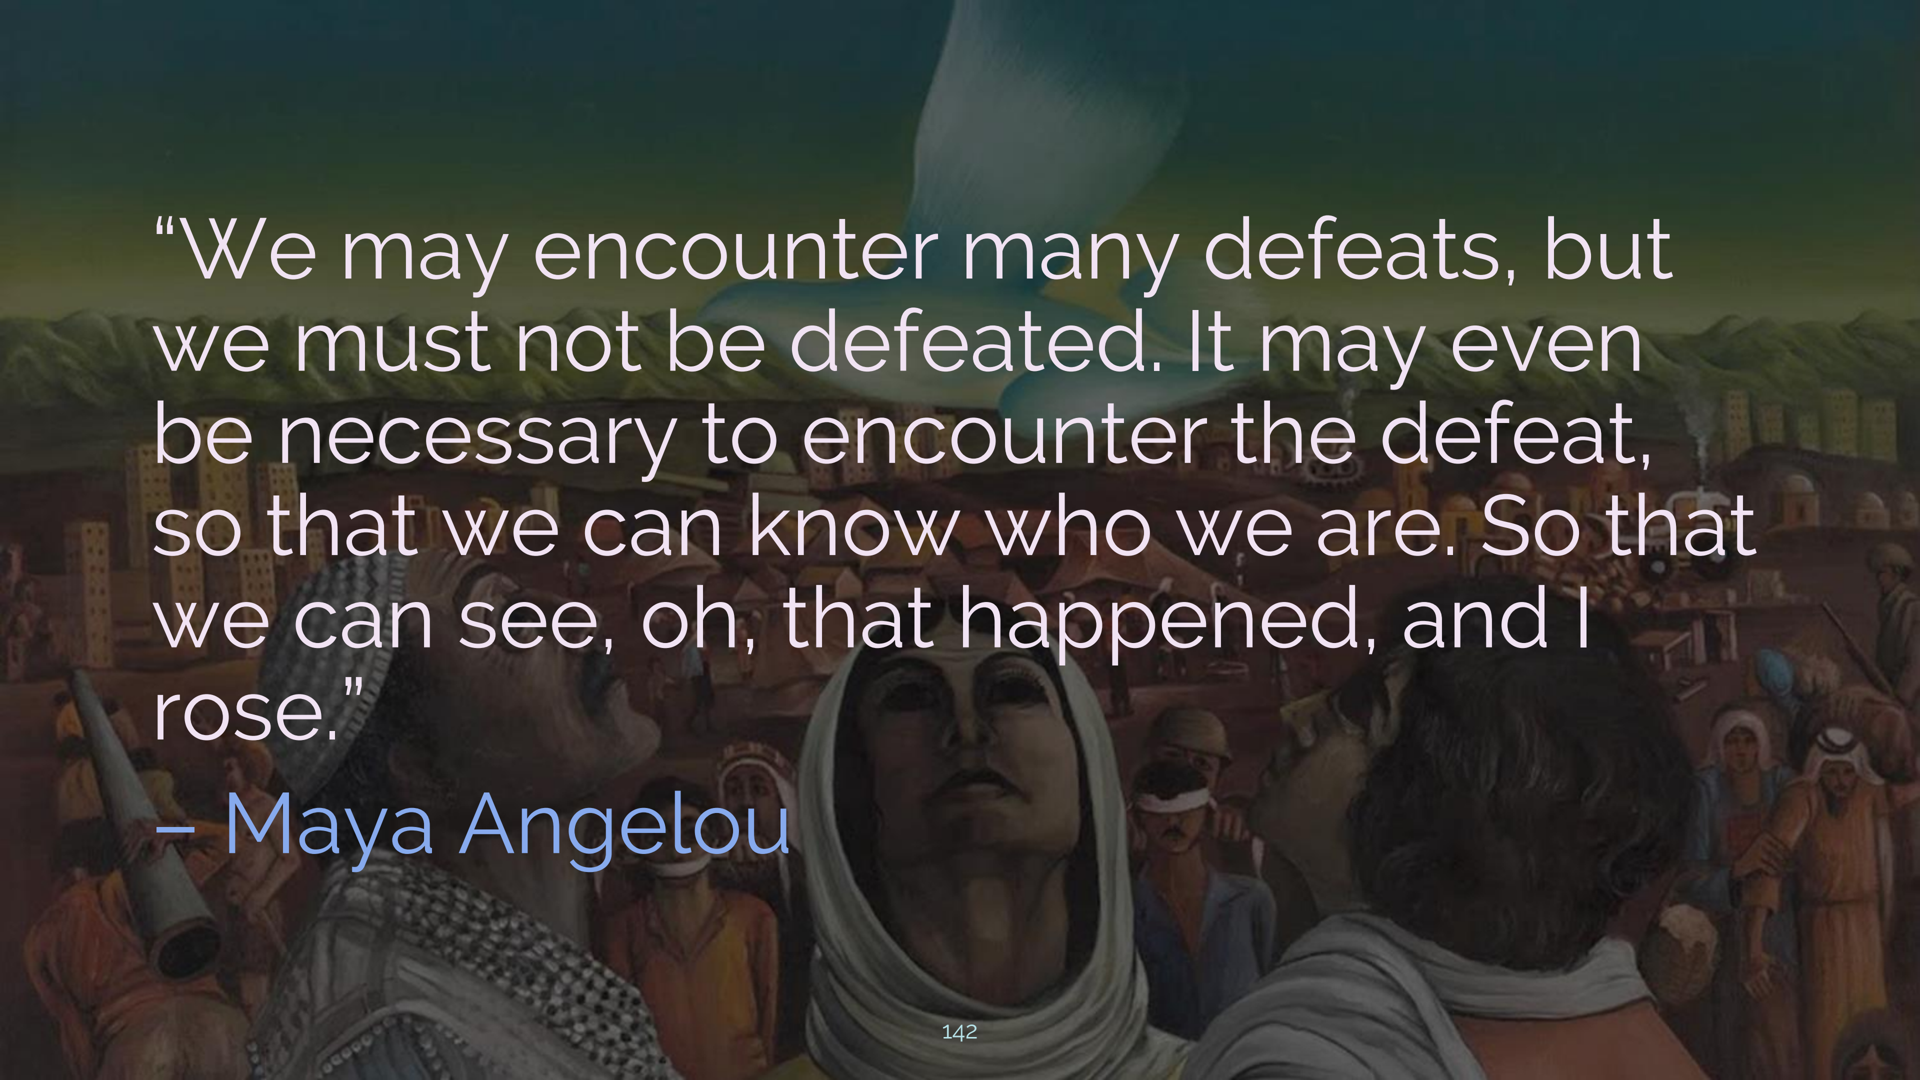

“We may encounter many defeats, but we must not be defeated. It may even be necessary to encounter the defeat, so that we can know who we are. So that we can see, oh, that happened, and I rose.”

– Maya Angelou

@swagitda_

/in/kellyshortridge

kelly@greywire.net

Suggested Reading

- "The evolution of error: Error management, cognitive constraints, and adaptive decision-making biases." Johnson, D., et al.
- "Hindsight bias impedes learning." Mahdavi, S., & Rahimian, M. A.
- "Outcome bias in decision evaluation." Baron, J., & Hershey, J. C.
- "Human error." Reason, J.
- "Behind human error." Woods, D., et al.
- "People or systems? To blame is human. The fix is to engineer." Holden, R.J.
- "Understanding adverse events: a human factors framework." Henriksen, K., et al.
- "Engineering a safer world: Systems thinking applied to safety." Leveson, N.
- "'Going solid': a model of system dynamics and consequences for patient safety." Cook, R., Rasmussen, J.
- "Choice Architecture." Thaler, R. H., Sunstein, C.R., Balz, J.P.
- "Blameless PostMortems and a Just Culture." Allspaw, J.