A Dangerous Folly: Why Individual Attack Prediction Can’t Be Our Goal

Kelly Shortridge (@swagitda_)

Art into Science 2018

Hi, I’m Kelly

SecurityScorecard

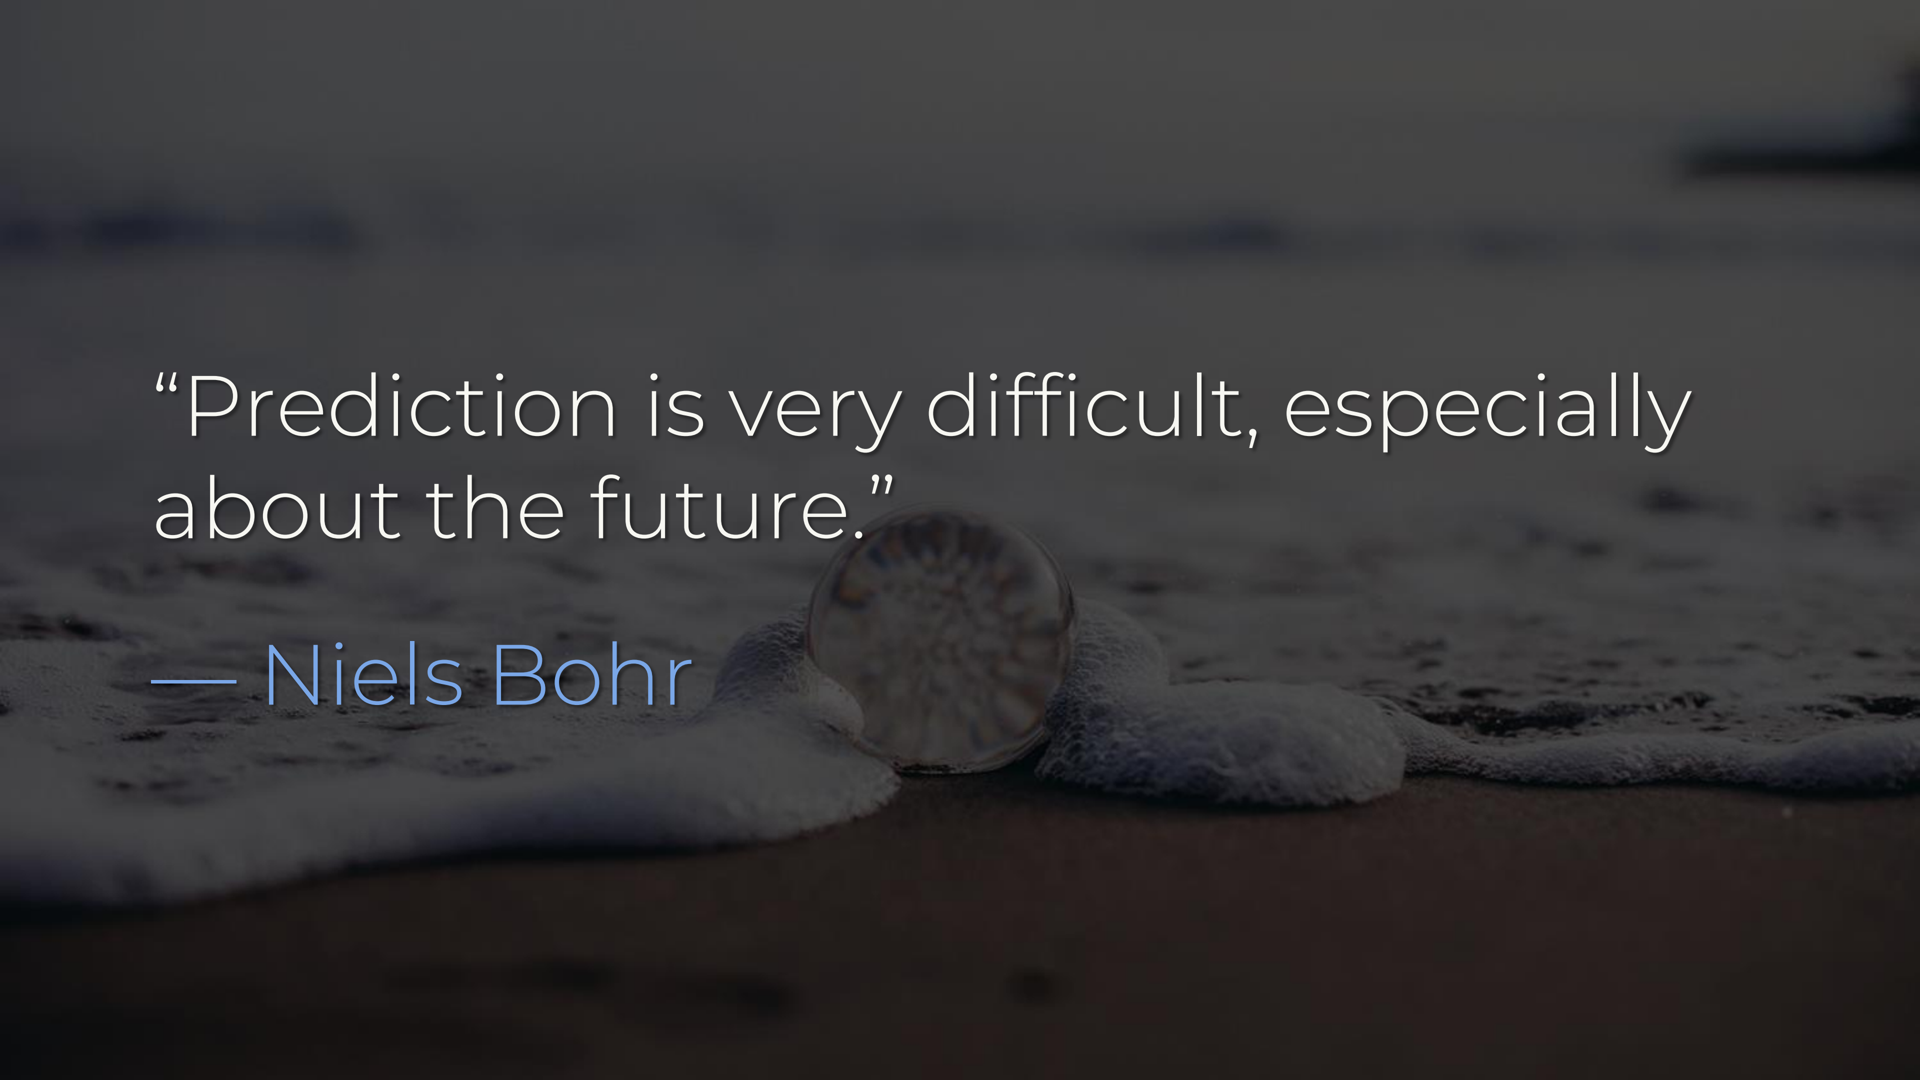

> “Prediction is very difficult, especially about the future.”
>
> — Niels Bohr

Problem: prediction is a sexy problem

Designing building codes is not

Act 1:

Why is everyone hyped on prediction & what methods do they propose?

Act 2:

What lessons exist from other areas & what should we do instead?

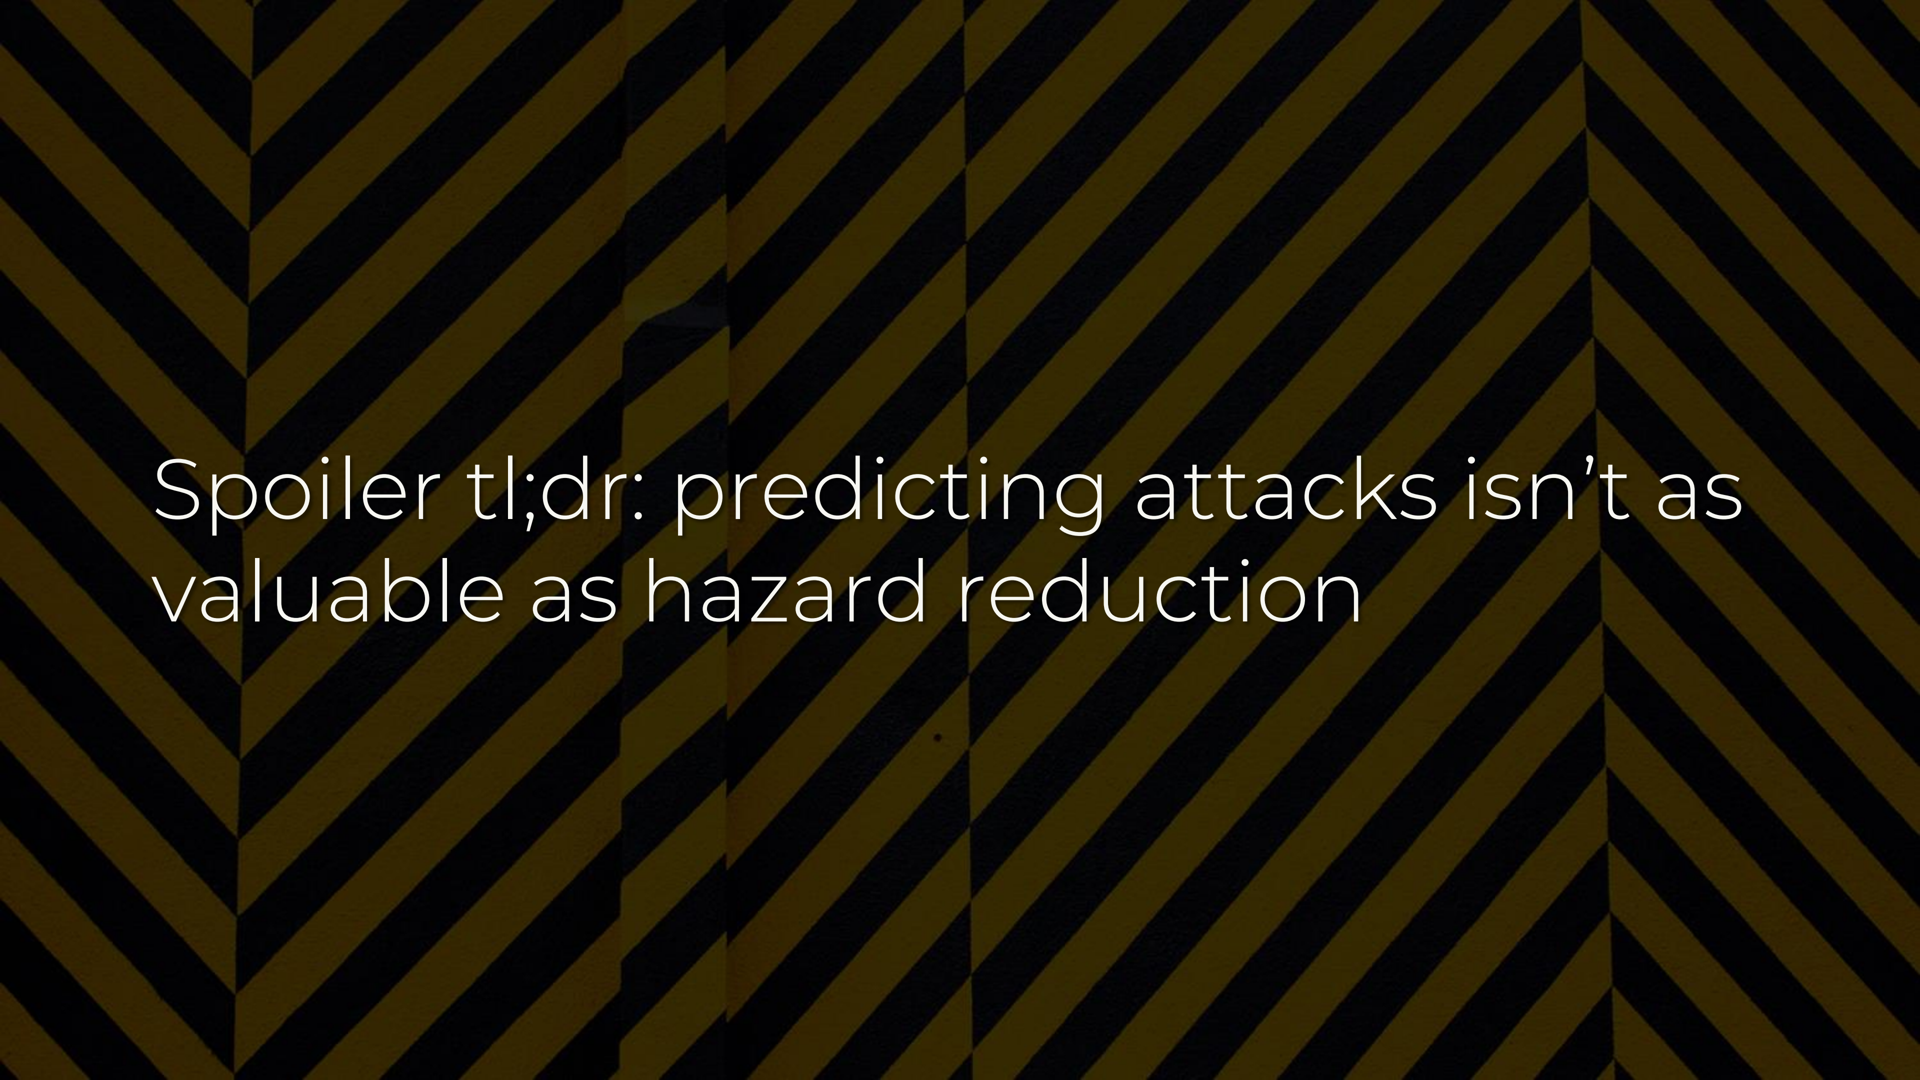

Spoiler tl;dr: predicting attacks isn’t as valuable as hazard reduction

Act 1

Why the interest in attack prediction?

Fundamentally, uncertainty feels bad

Ambiguity of potential future threats fuels stress & anxiety

Predictable negative events are less stressful than uncertainty

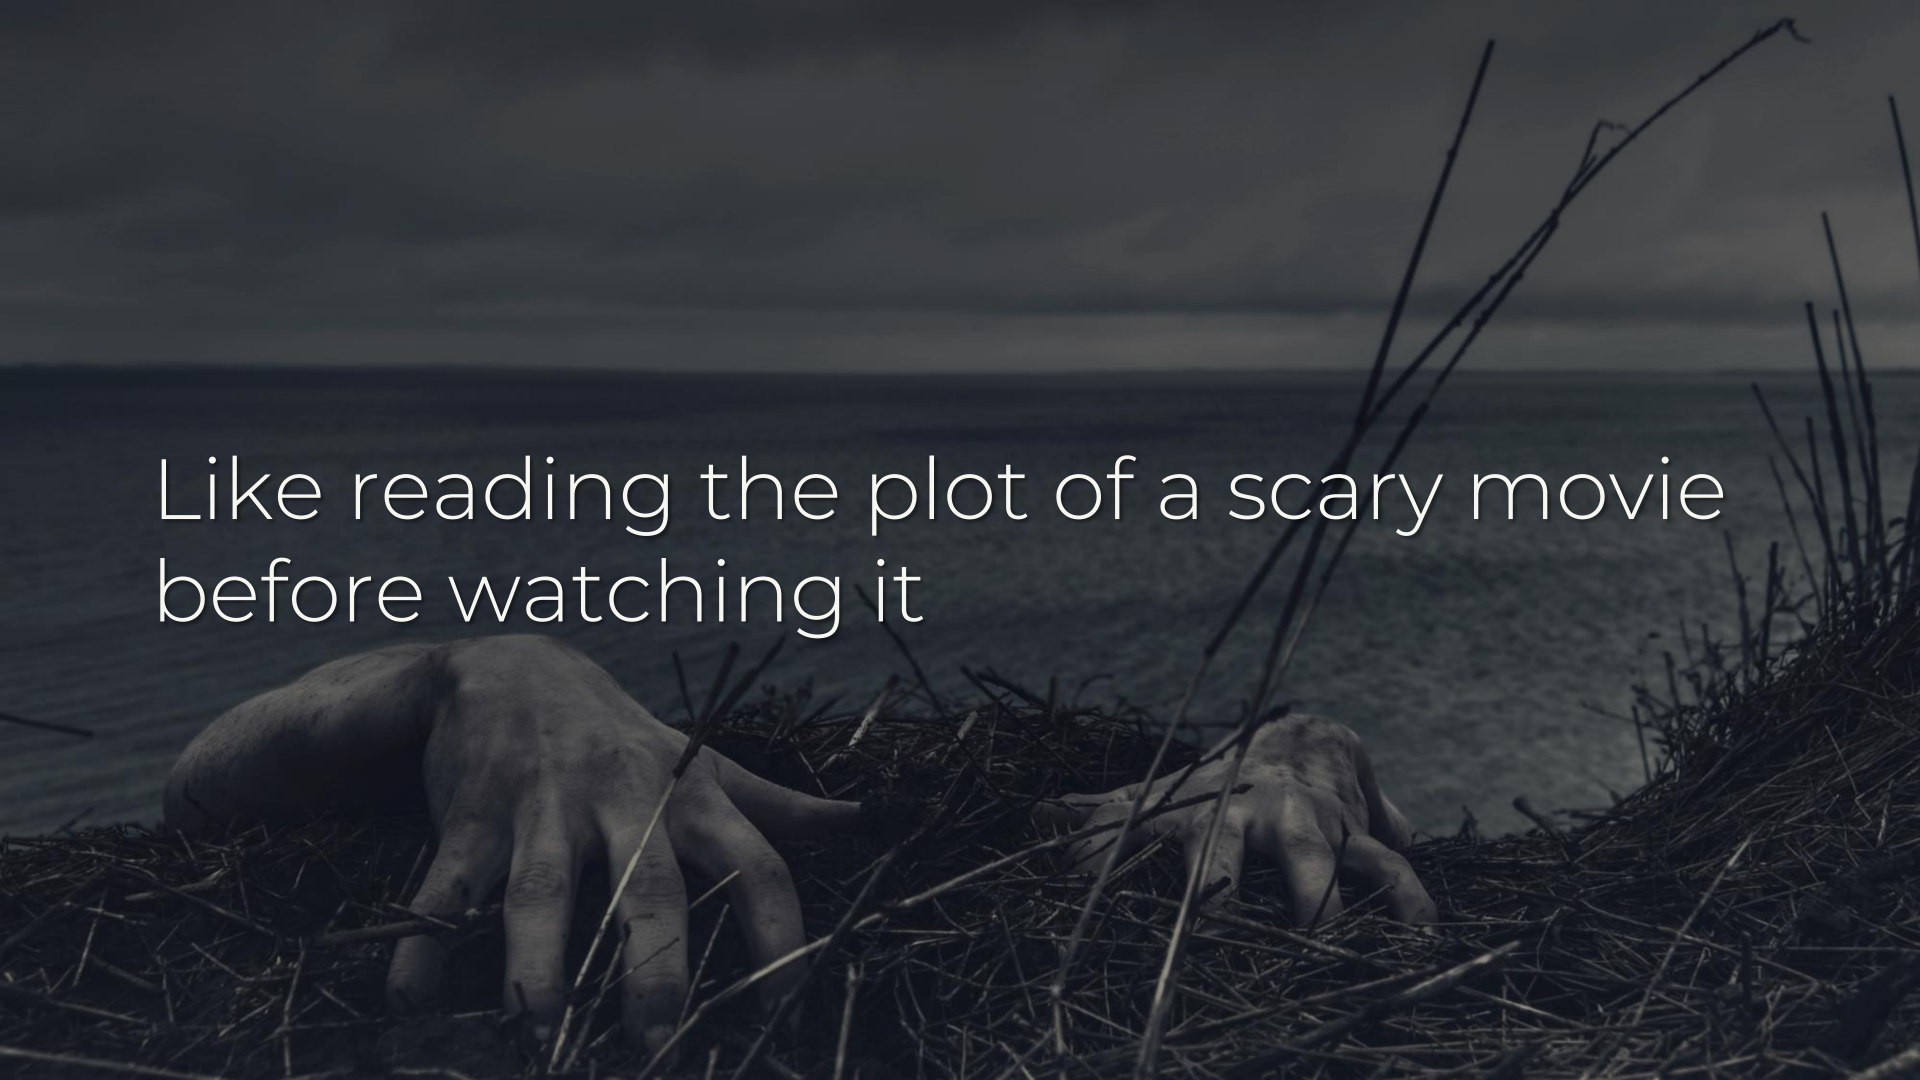

Like reading the plot of a scary movie before watching it

Unfortunately, predictions can give a false sense of security

There are some “unemotional” reasons used to justify prediction, too

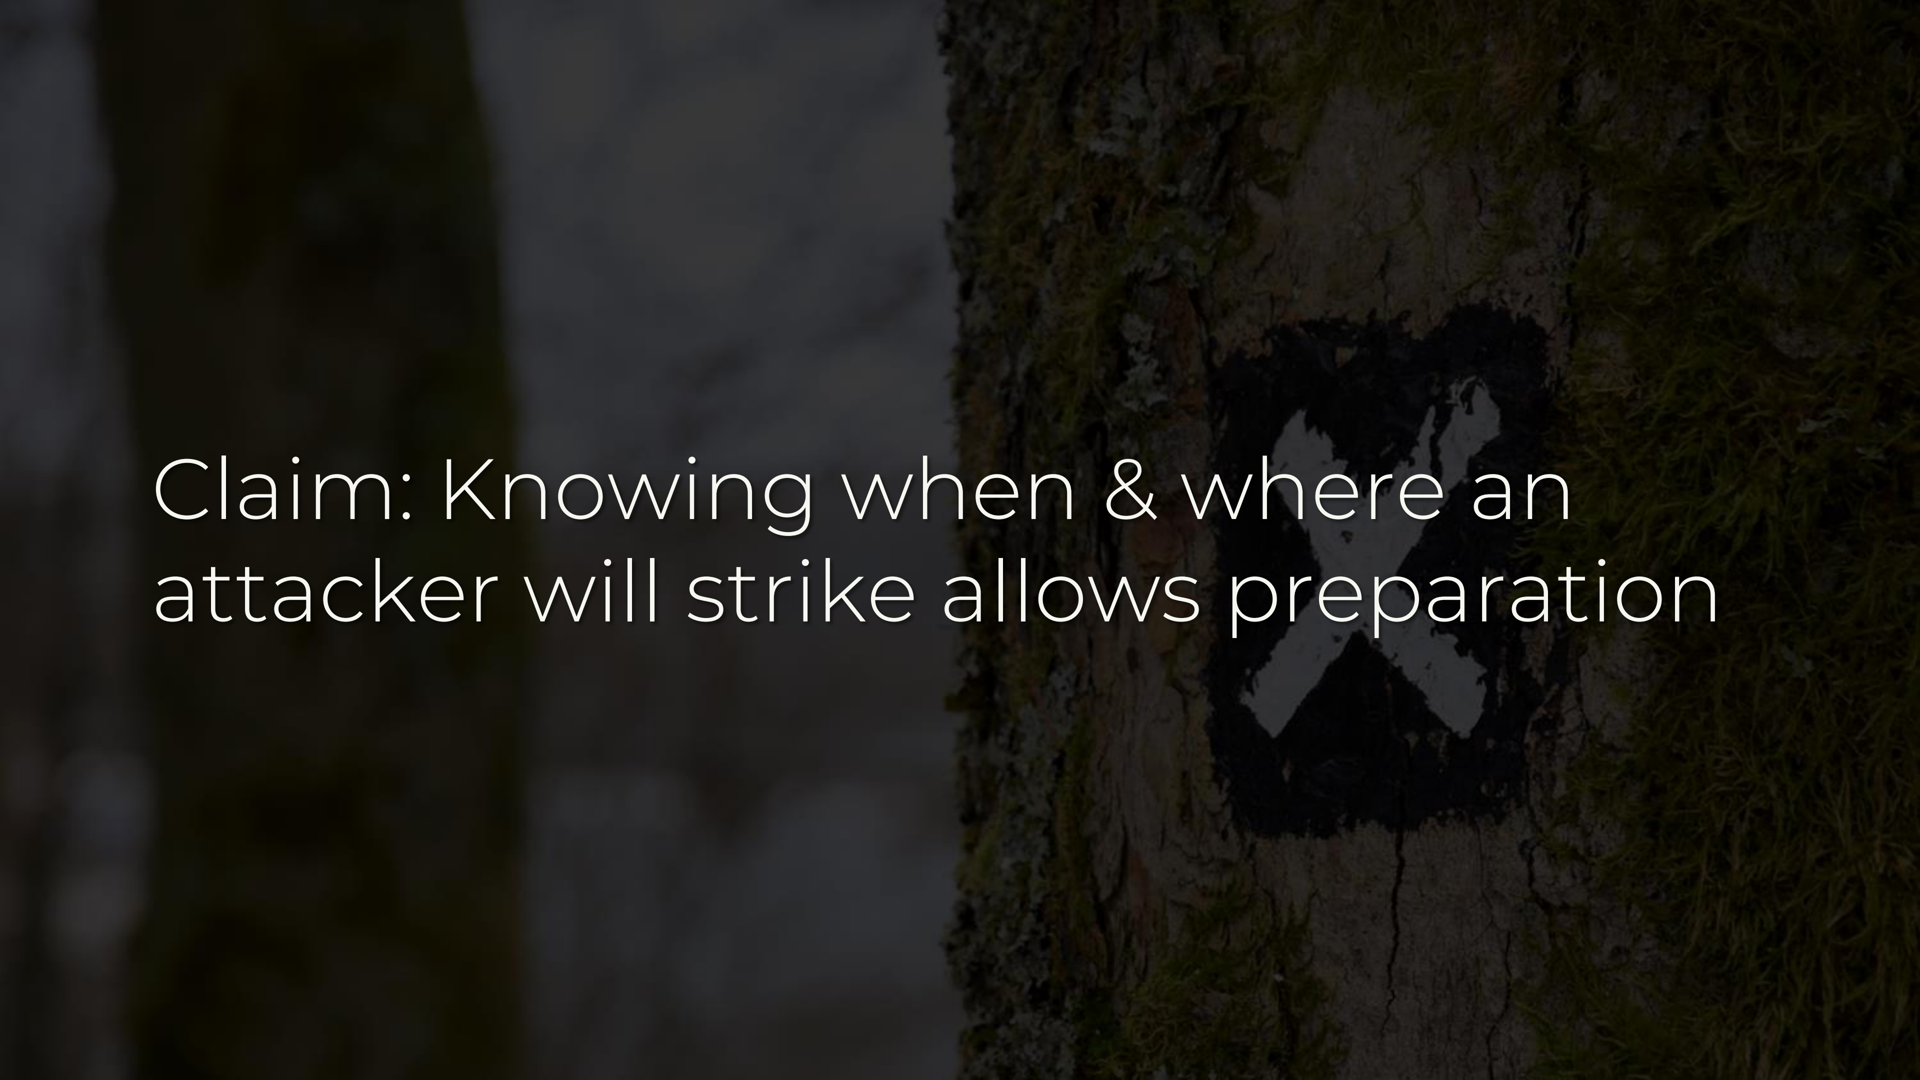

Claim: Knowing when & where an attacker will strike allows preparation

Claim: Knowing the attacker’s next move helps with resource allocation

Tacit reason: precogs are cool

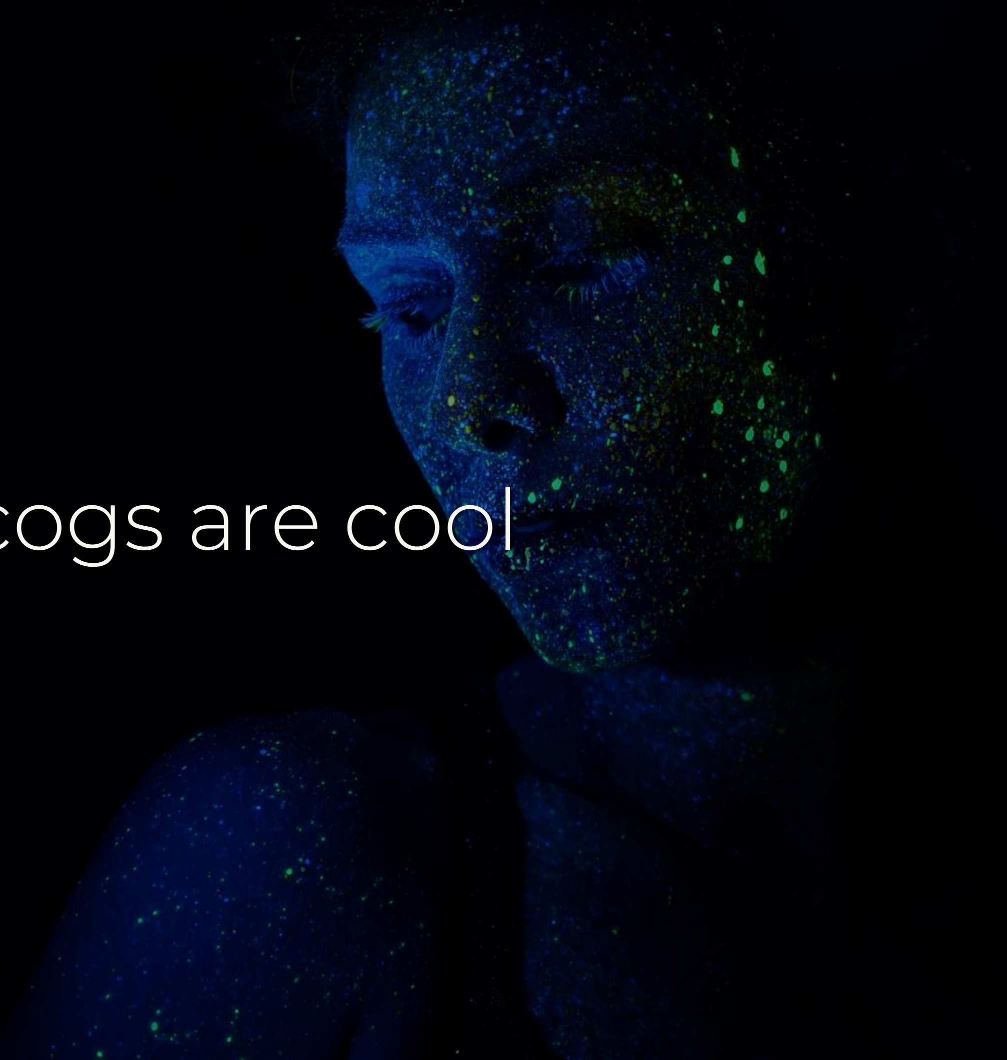

But do the goals align with the methods actually being proposed?

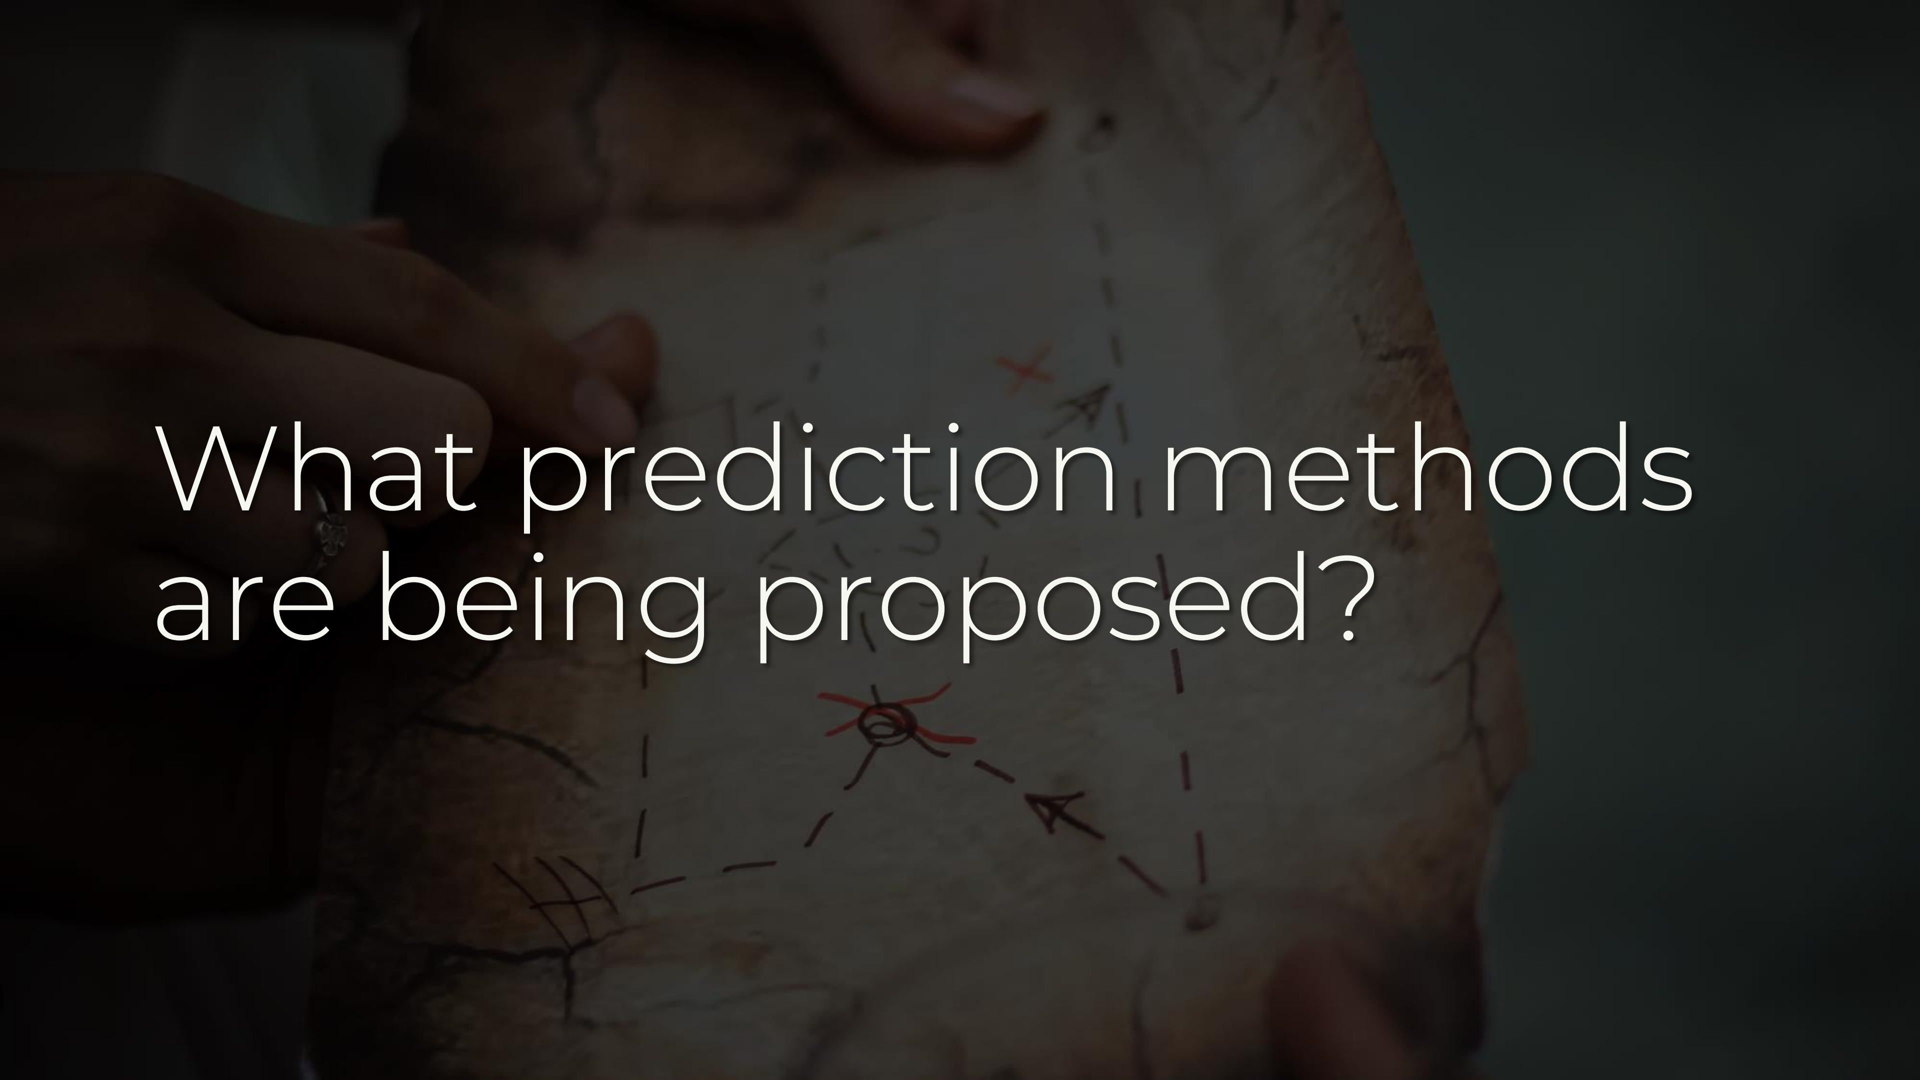

What prediction methods are being proposed?

General theme: predict future attacks from past & current attack behavior

Why now? Sufficient storage, processing power, & we math better

mysheen lerning mysheene lerning
masheene lerning mashene learning
mashine learning machine learning

Common idea: unsupervised machine learning to avoid false negatives

AI – i.e. a magic black box of math

Deep learning threat prediction using actionable behavioral analytics…

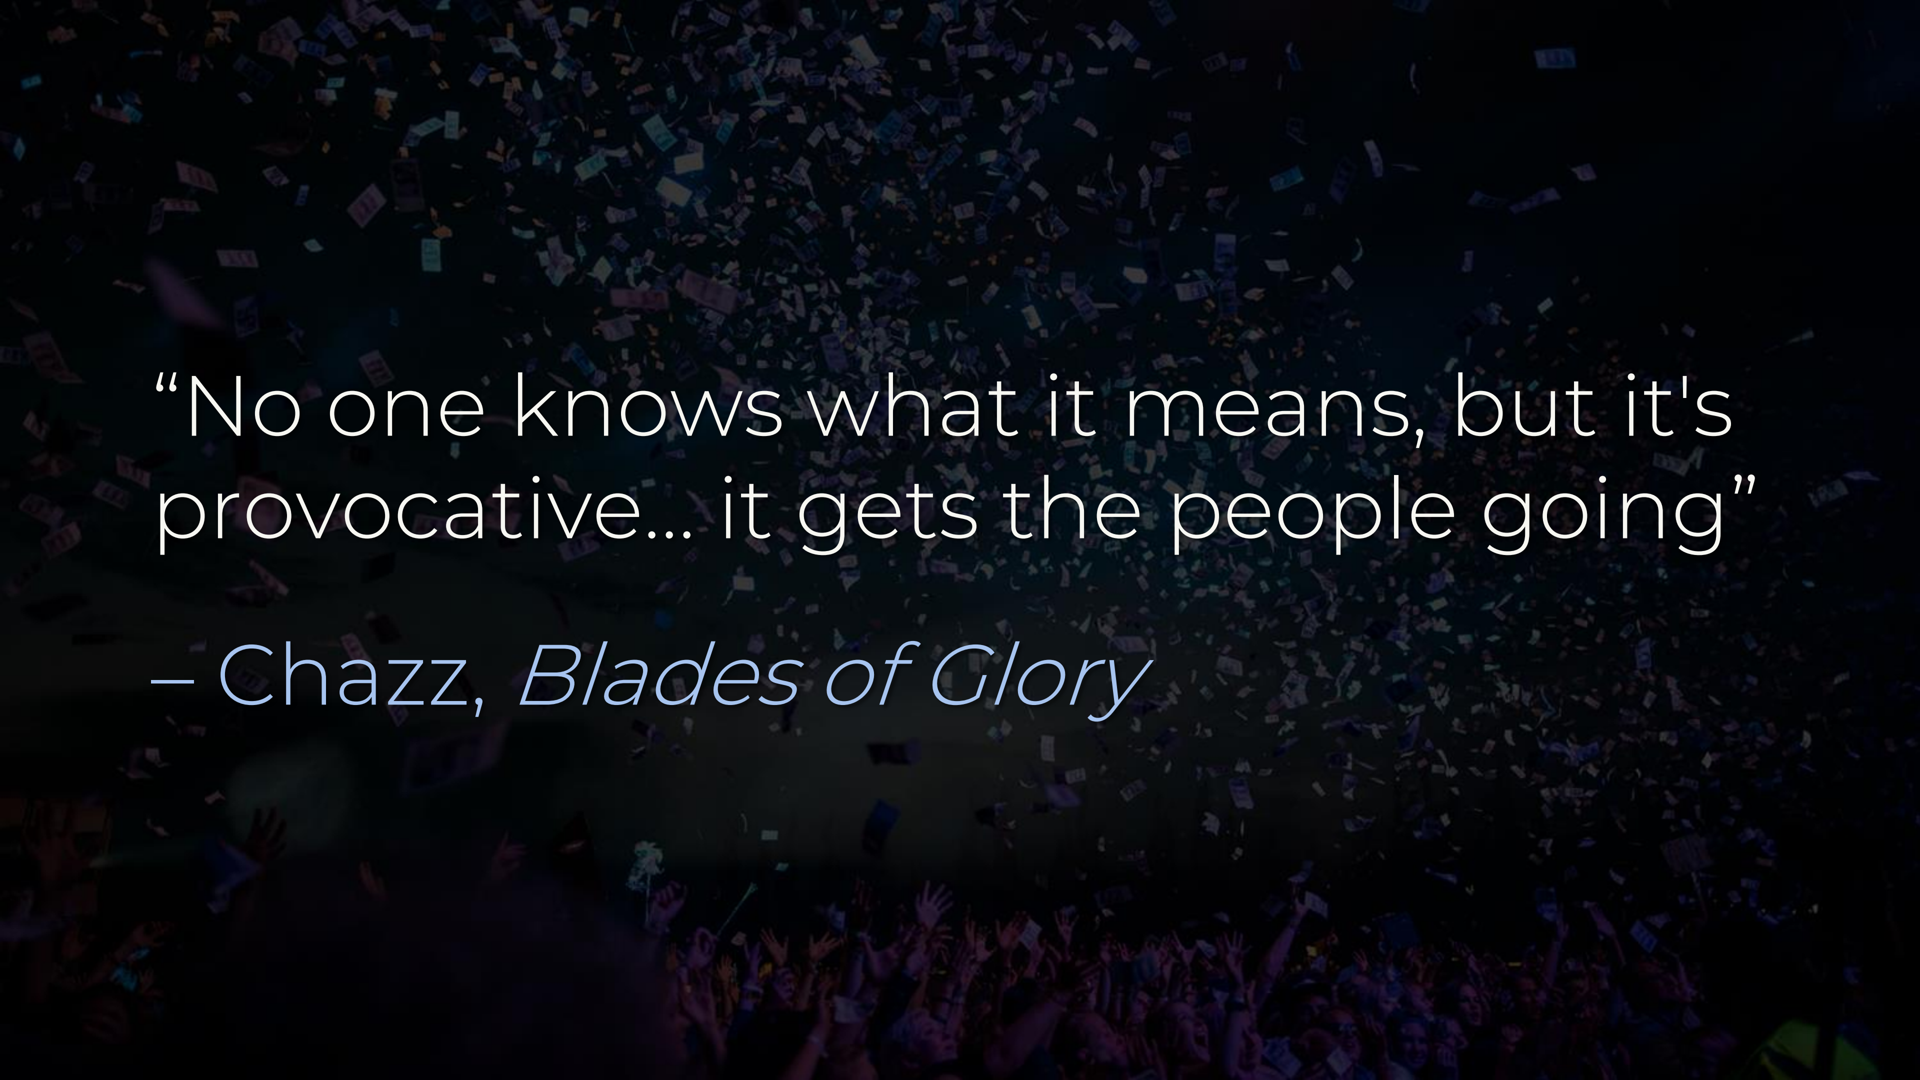

“No one knows what it means, but it's provocative… it gets the people going”

– Chazz, *Blades of Glory*

Caveat: behavioral analytics for detection already plagued by FPs

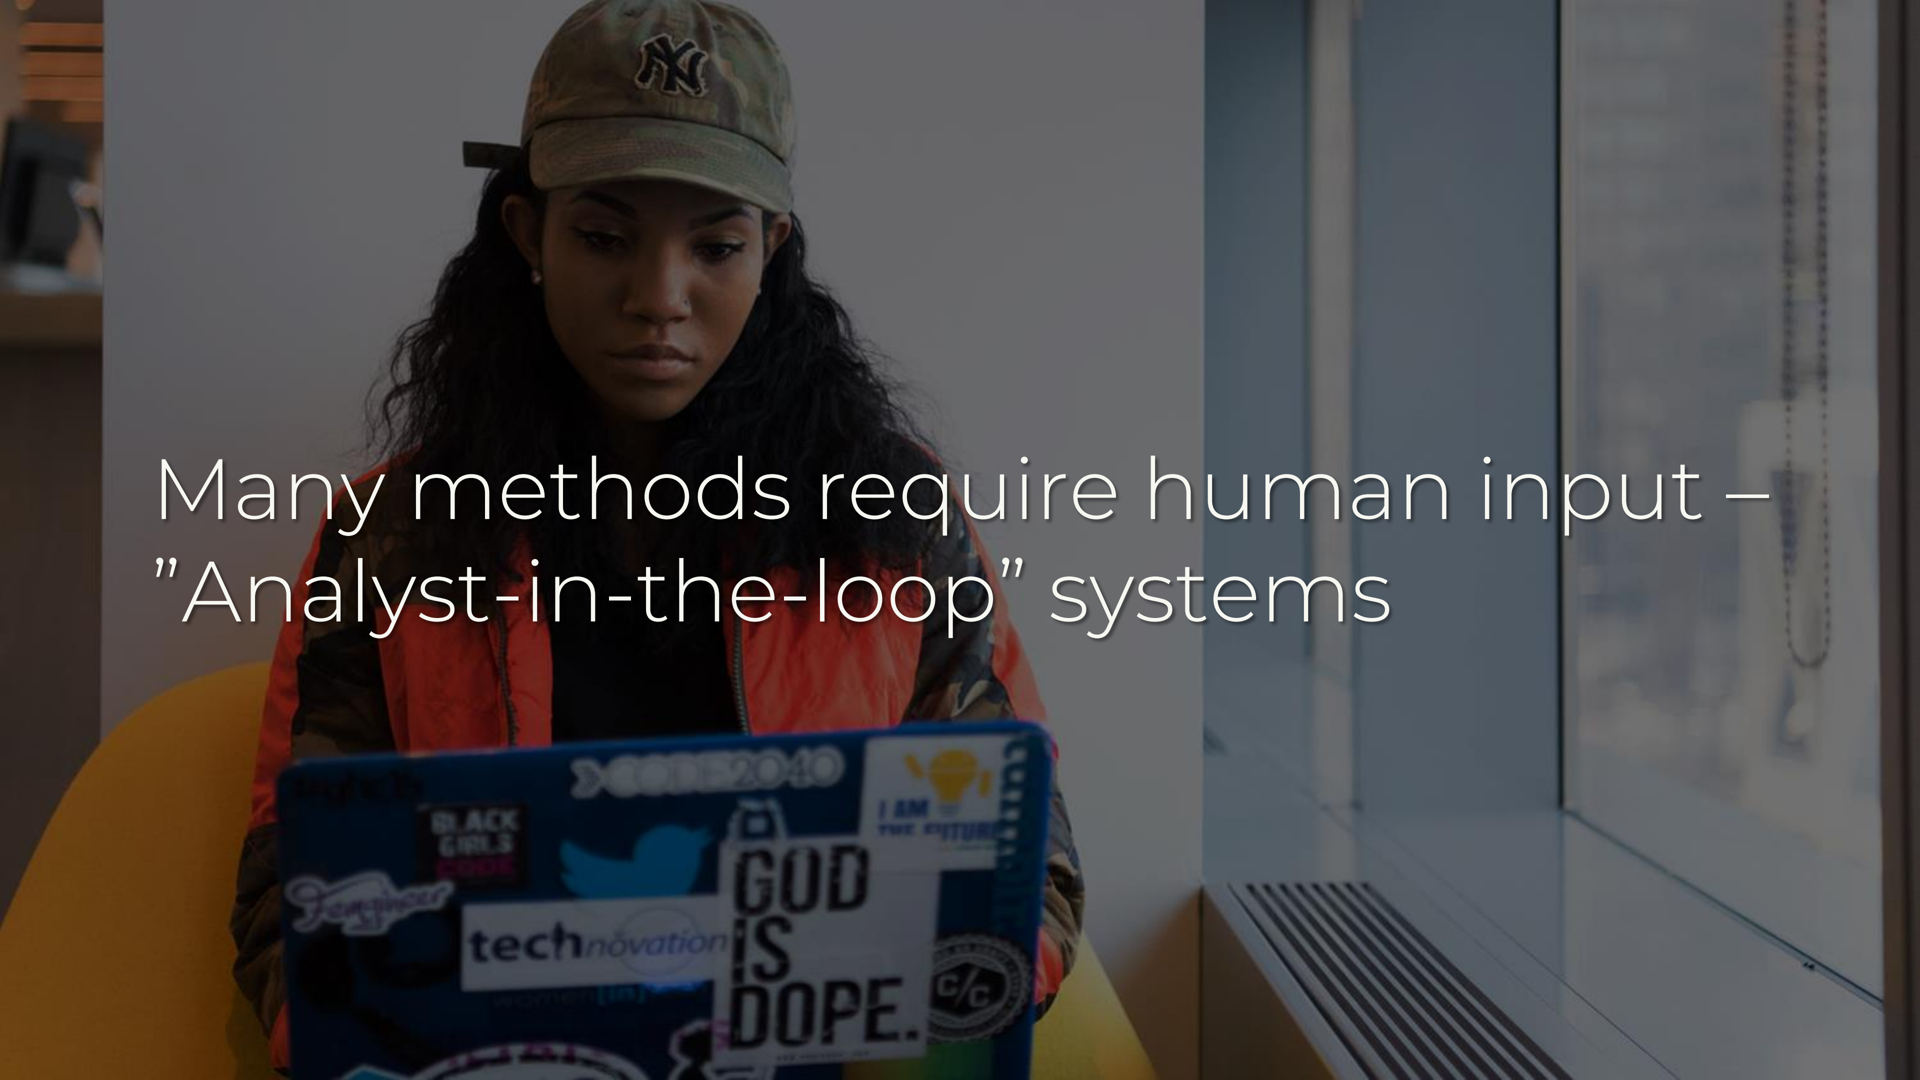

Many methods require human input – ”Analyst-in-the-loop” systems

AI^2 : fuses 3 unsupervised-learning methods & shows top events to analysts for them to label

Caveat: still focused on detection – calling it “prediction” is a stretch

(Also, the paper reads like an ad for IBM Watson & QRadar…)

Idea: use attacker TTPs to train your data sci models

Caveat: it’s really hard to attribute TTPs, let alone collect them

Reallocating resources on-the-fly based on predictions? Good luck…

Idea: social data analysis (using OSINT) to predict data breaches

Caveat: does “news” always know about breaches before the org itself?

(Also, how is it predictive if they’re finding news post-breach?)

Cause analysis: what allowed the attack to happen?

Caveat: past performance is not an indicator of future performance

Plausible: detecting preparations helps stop attacks before they start

Caveat: how much does a predictive system add vs. using canaries?

Where does this leave us?

Barriers: FPs, attackers are quick to adapt behaviors, limited time/people

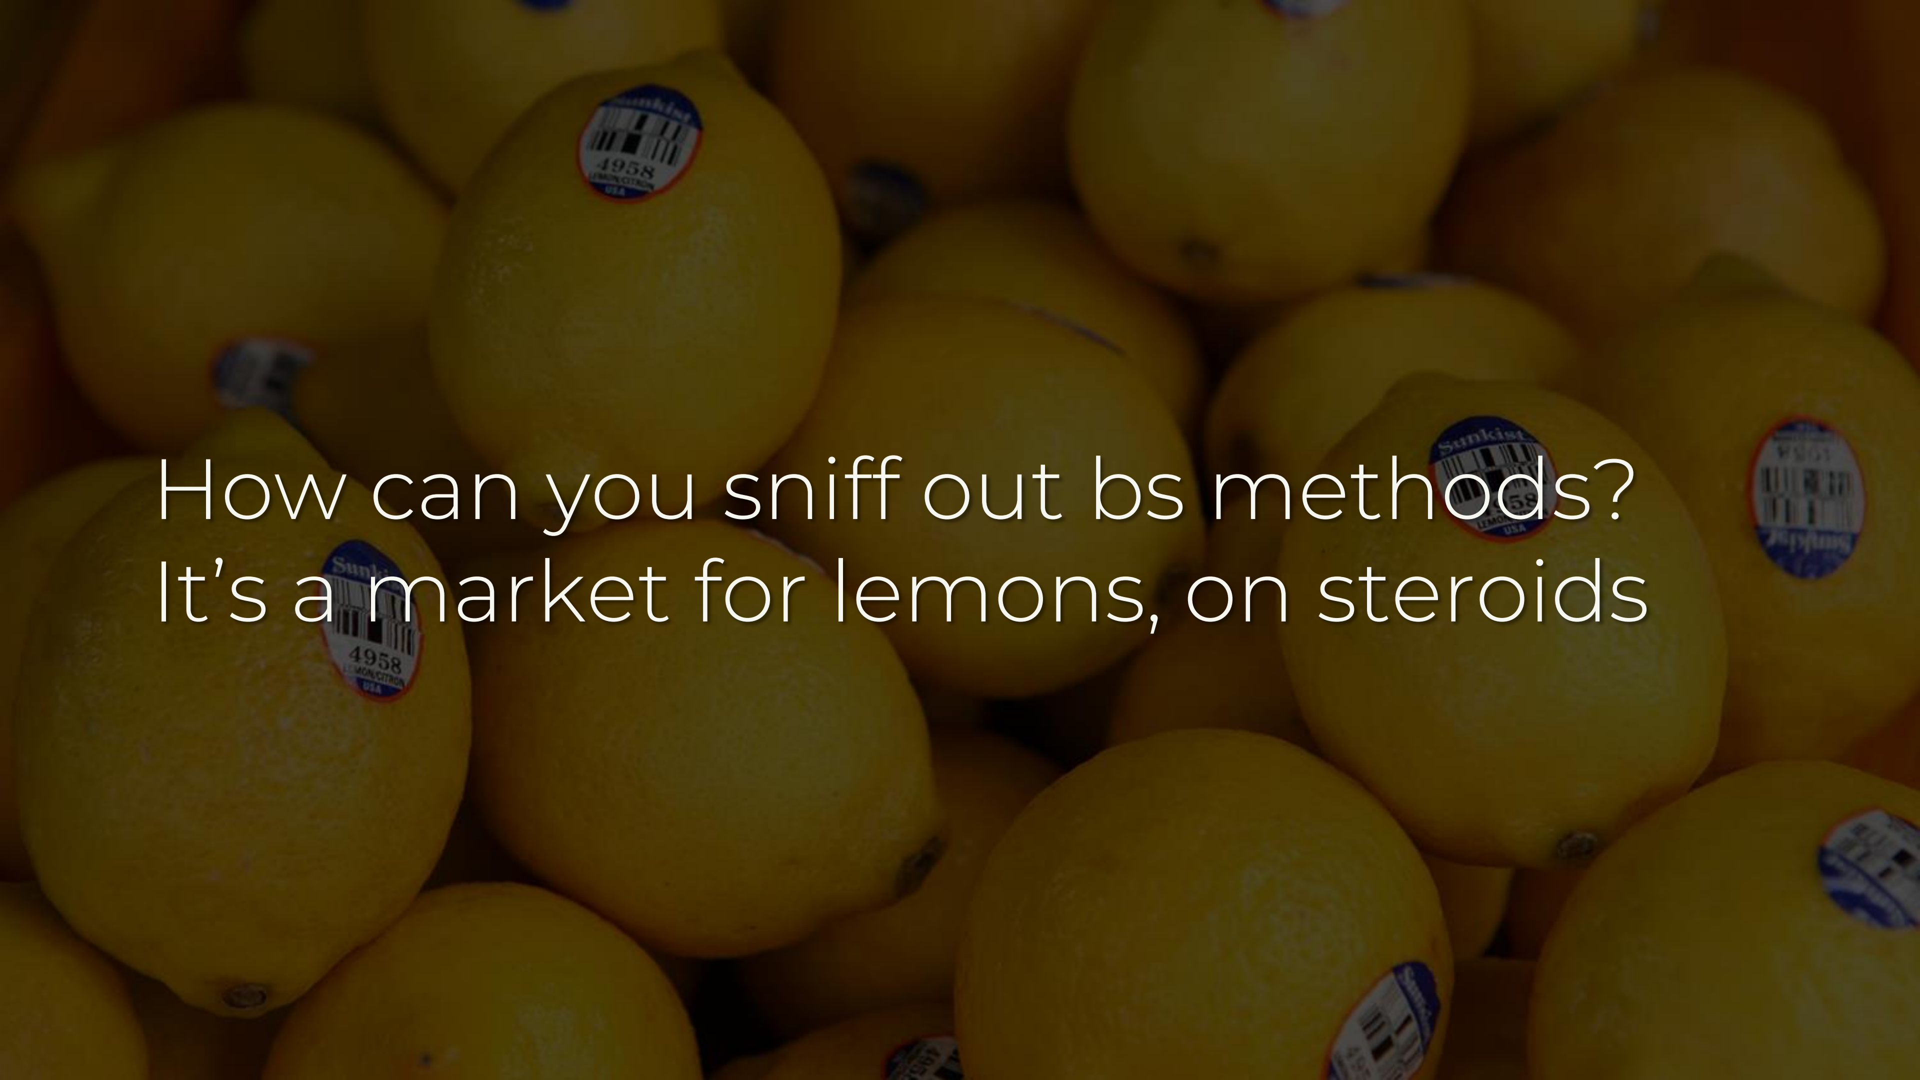

How can you sniff out bs methods?
It’s a market for lemons, on steroids

How do you know what to do next?
Predictive vs. prescriptive

Alternative proposal: prepare your prioritized assets for the (probabilistic) worst, ahead of attack…

Act 2

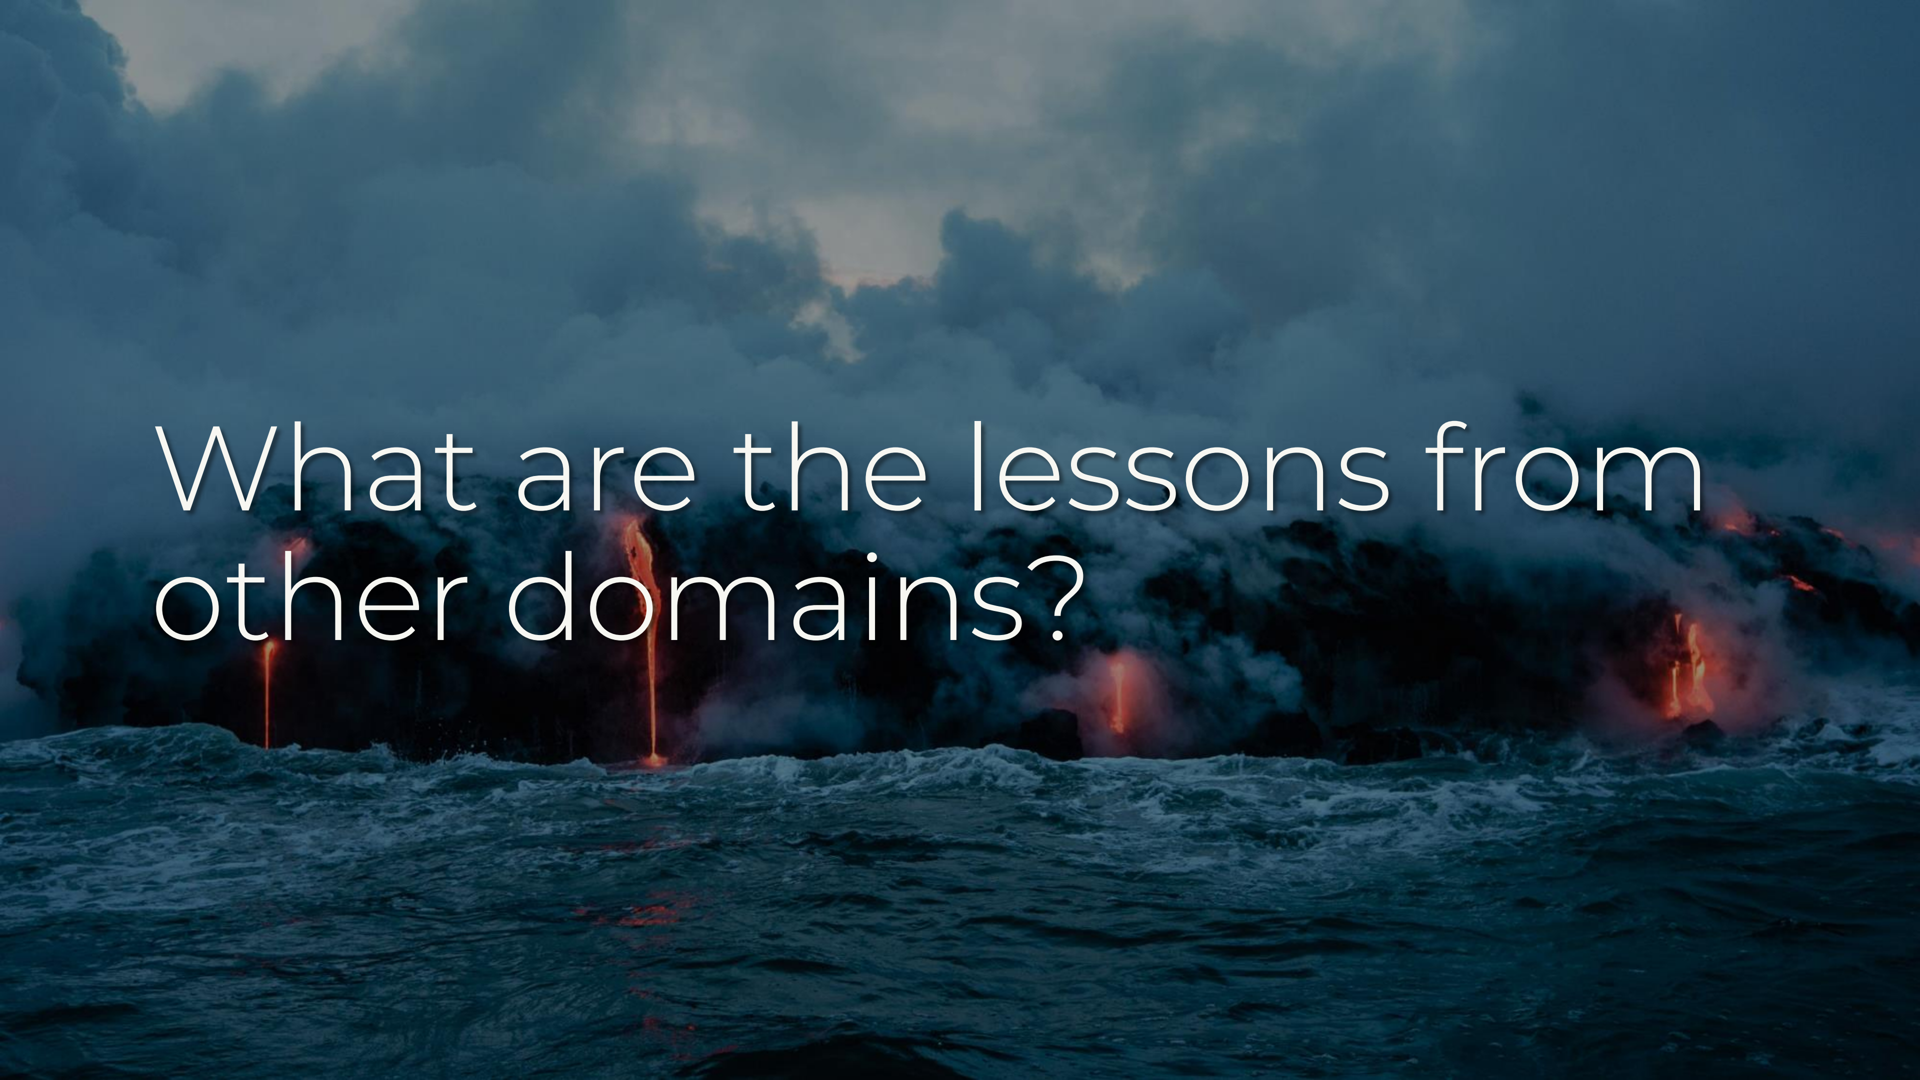

What are the lessons from other domains?

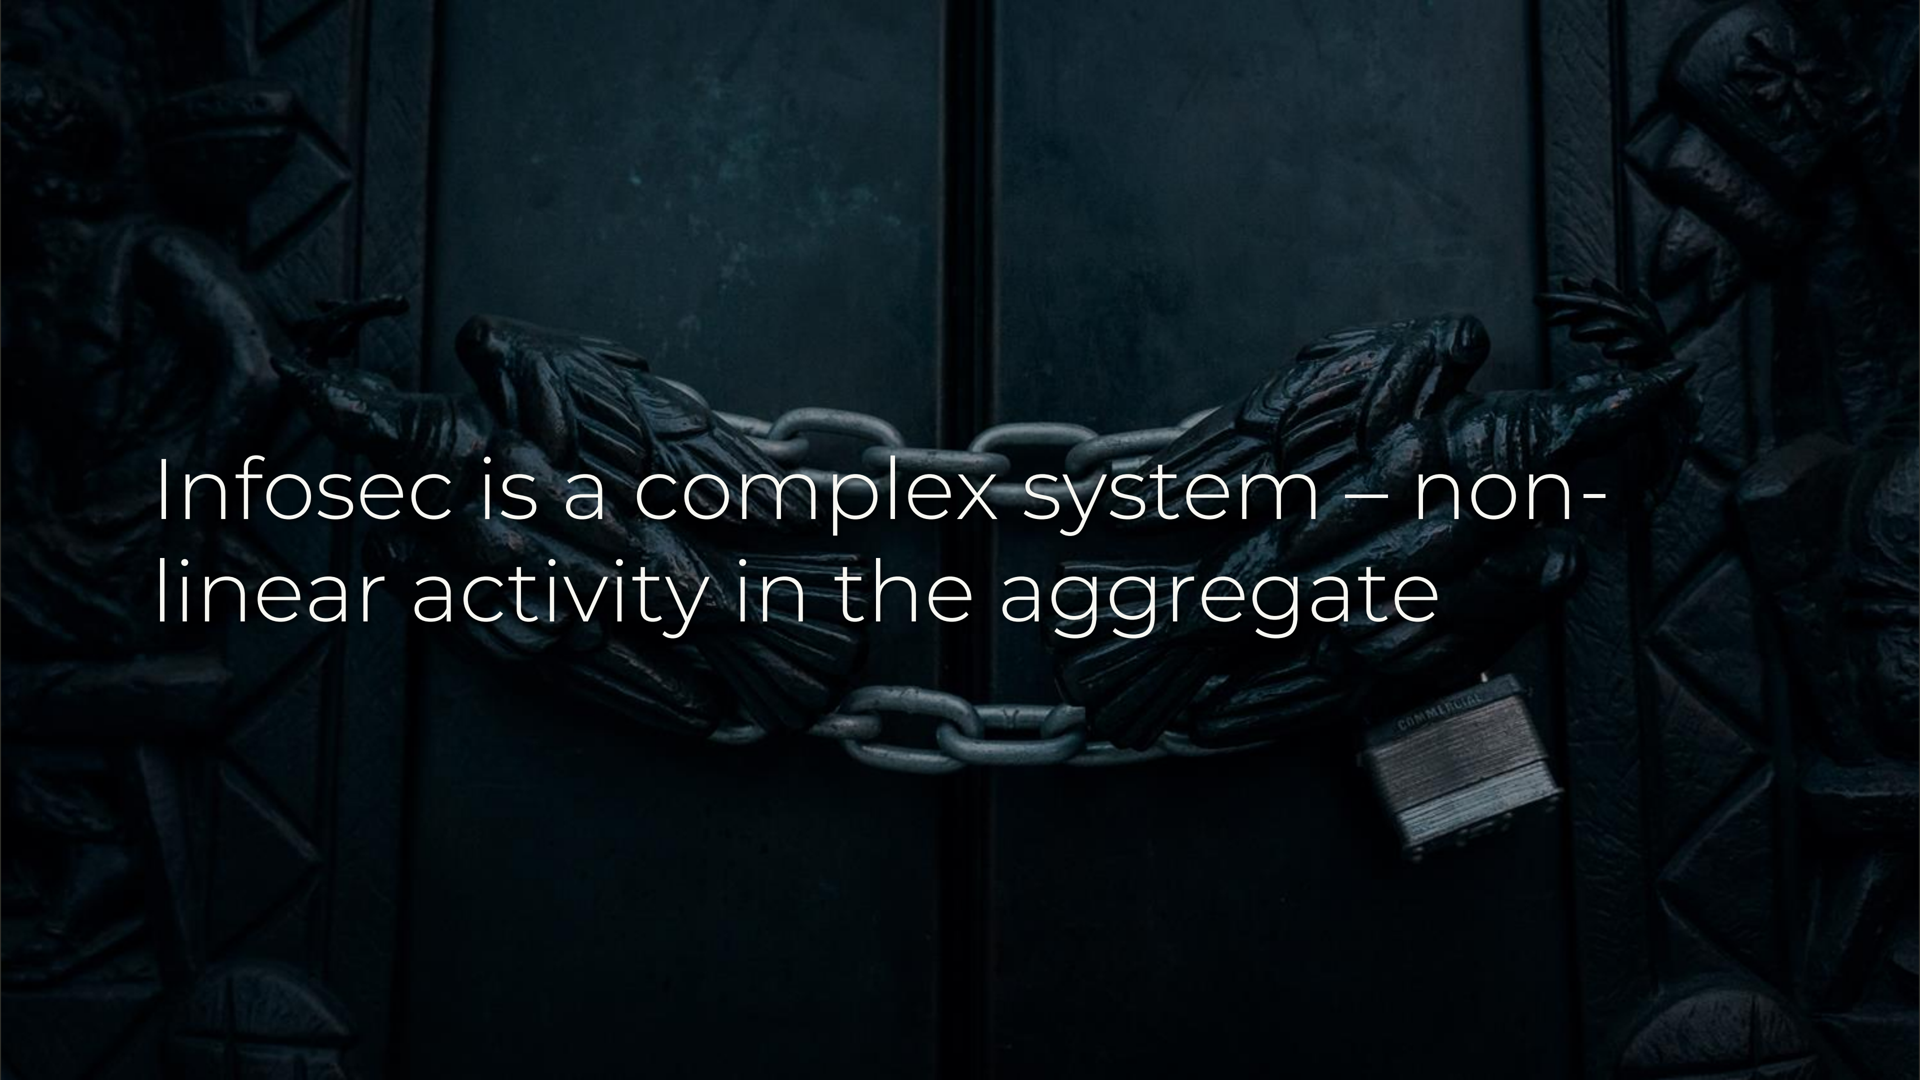

Infosec is a complex system – non-linear activity in the aggregate

Prediction of natural disasters = knowing time, location, and severity

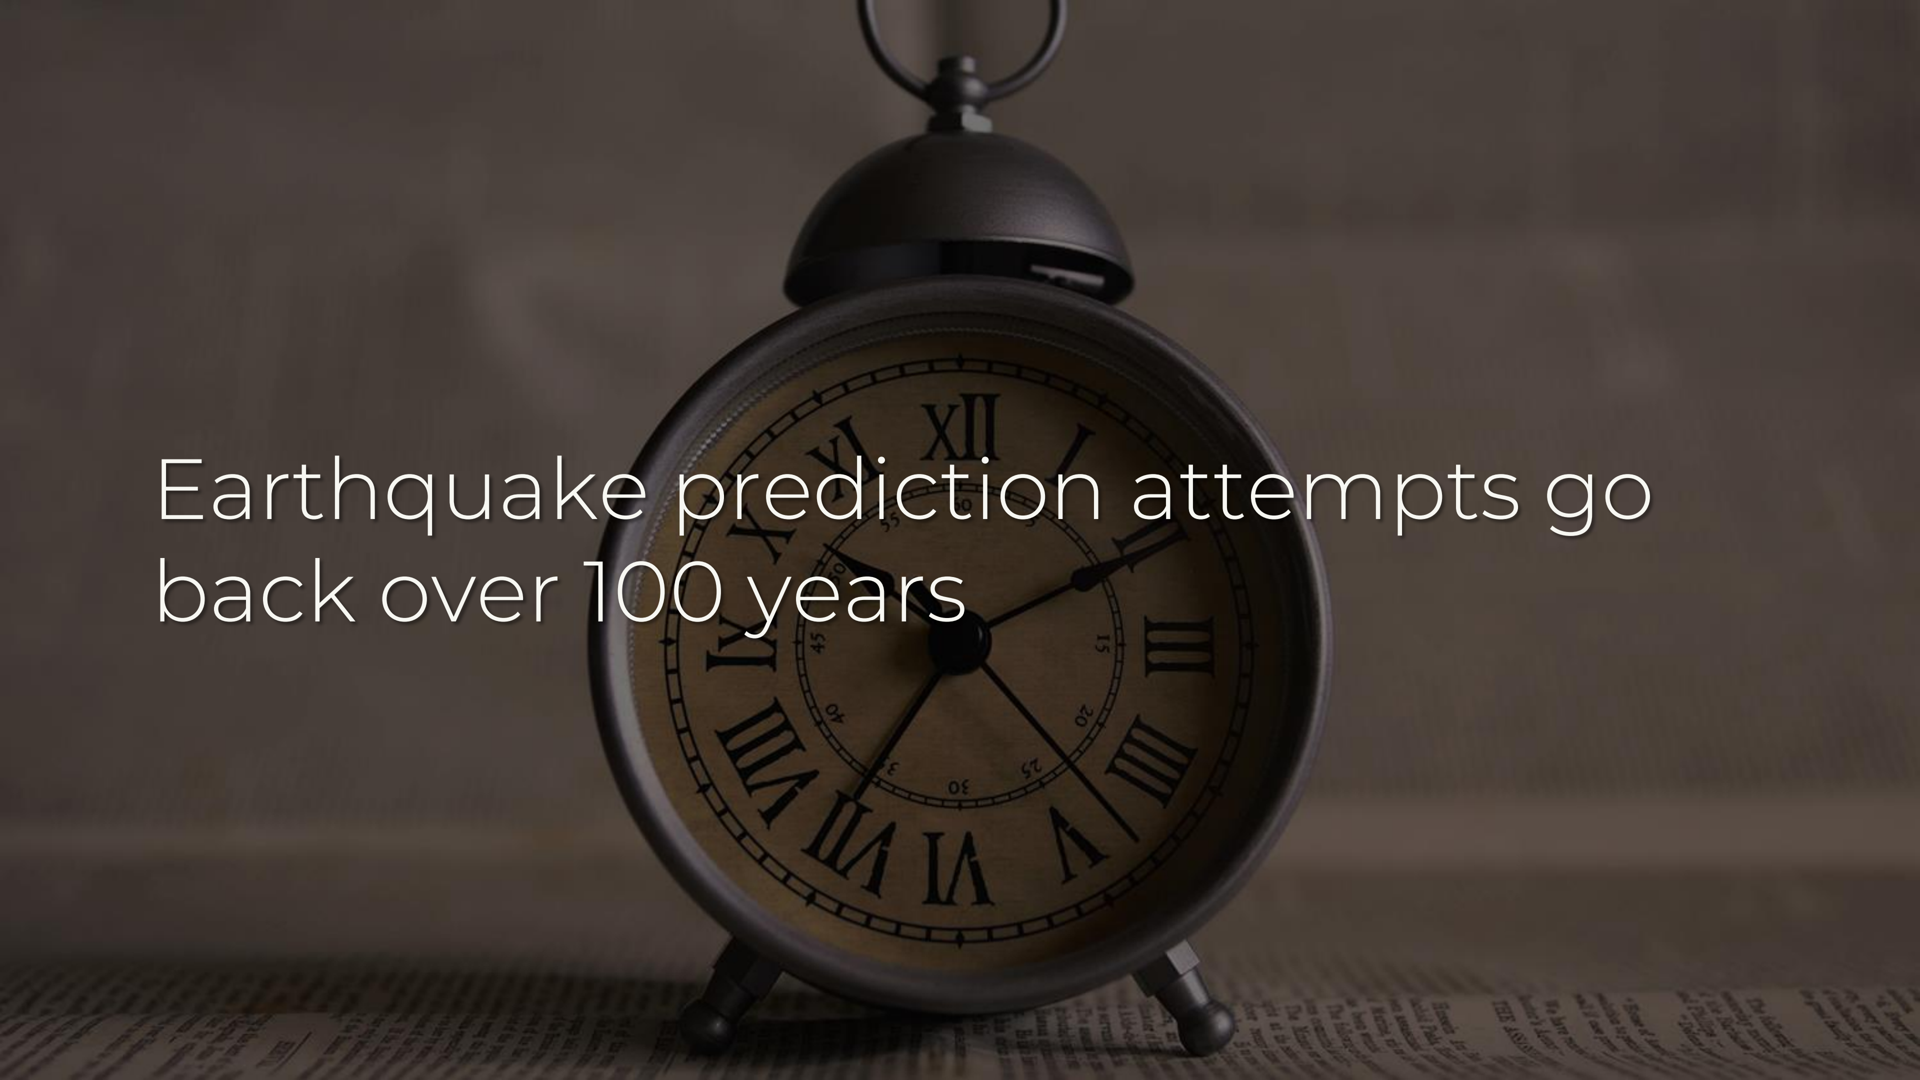

Earthquake prediction attempts go back over 100 years

1970s: success within the next 10 years

2000s: prediction is (probably) impossible – or at least far off

We still can’t predict earthquakes, despite tons of funding

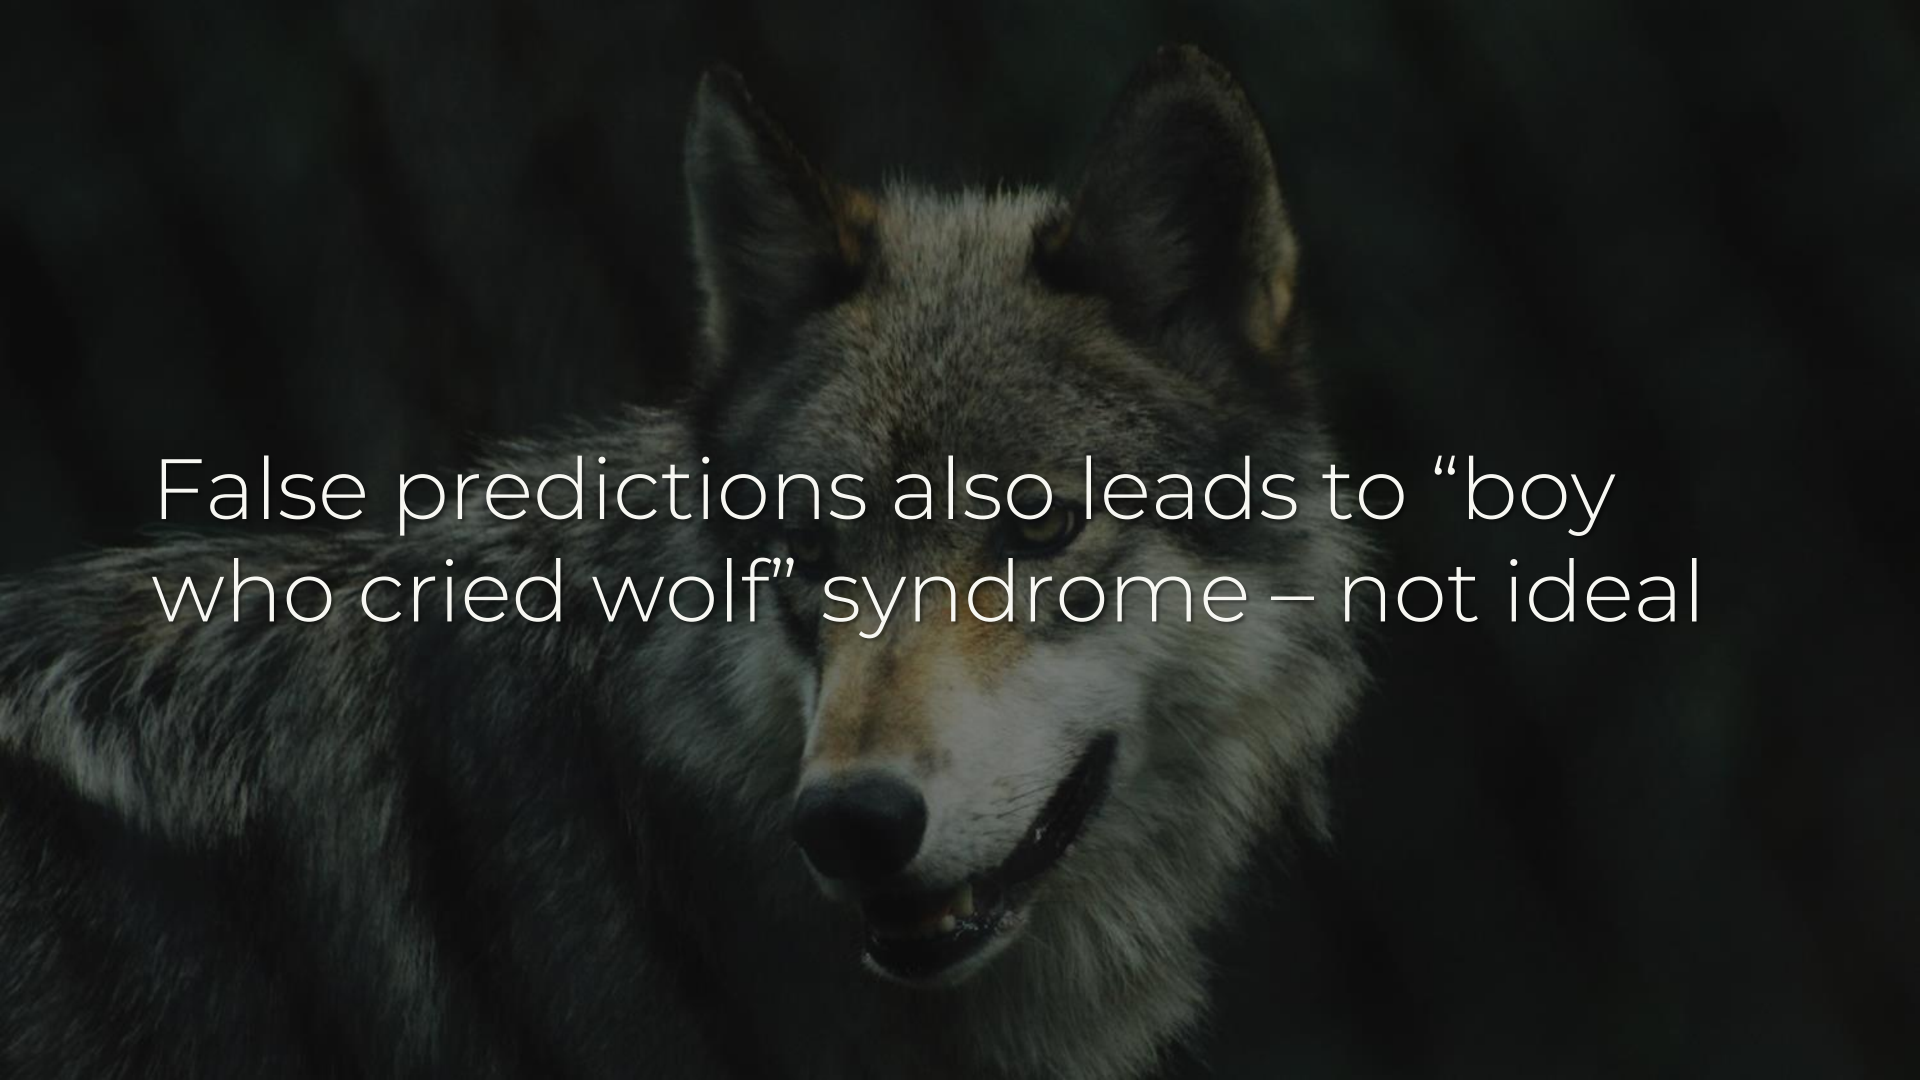

False predictions also leads to “boy who cried wolf” syndrome – not ideal

Earthquake **forecasting** vs. earthquake **prediction**

We know which areas are risky, but not where & when a quake will occur

This is enough info to inform us that we need to be prepared

Building codes: withstand effects & incur acceptable level of damage

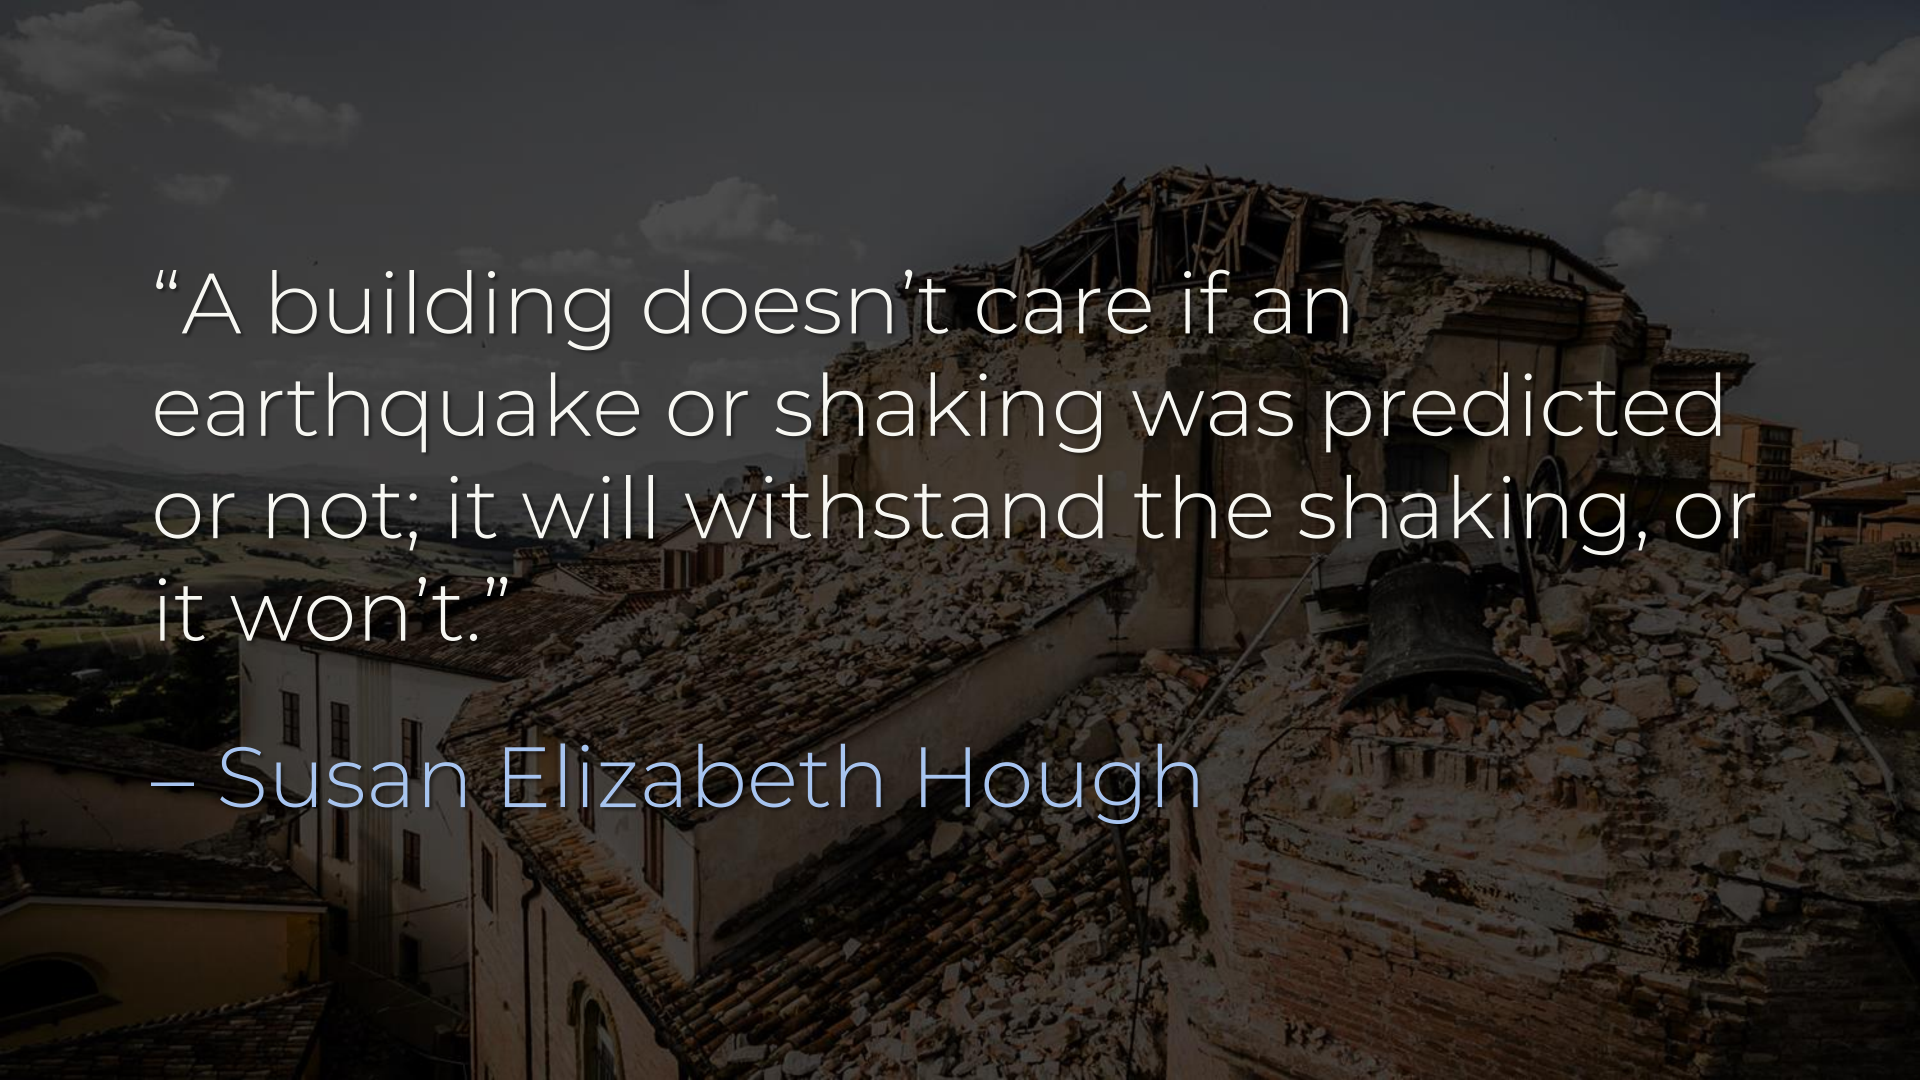

"A building doesn't care if an earthquake or shaking was predicted or not; it will withstand the shaking, or it won't."

– Susan Elizabeth Hough

More valuable: reducing vulns, risk assessment, understanding impacts

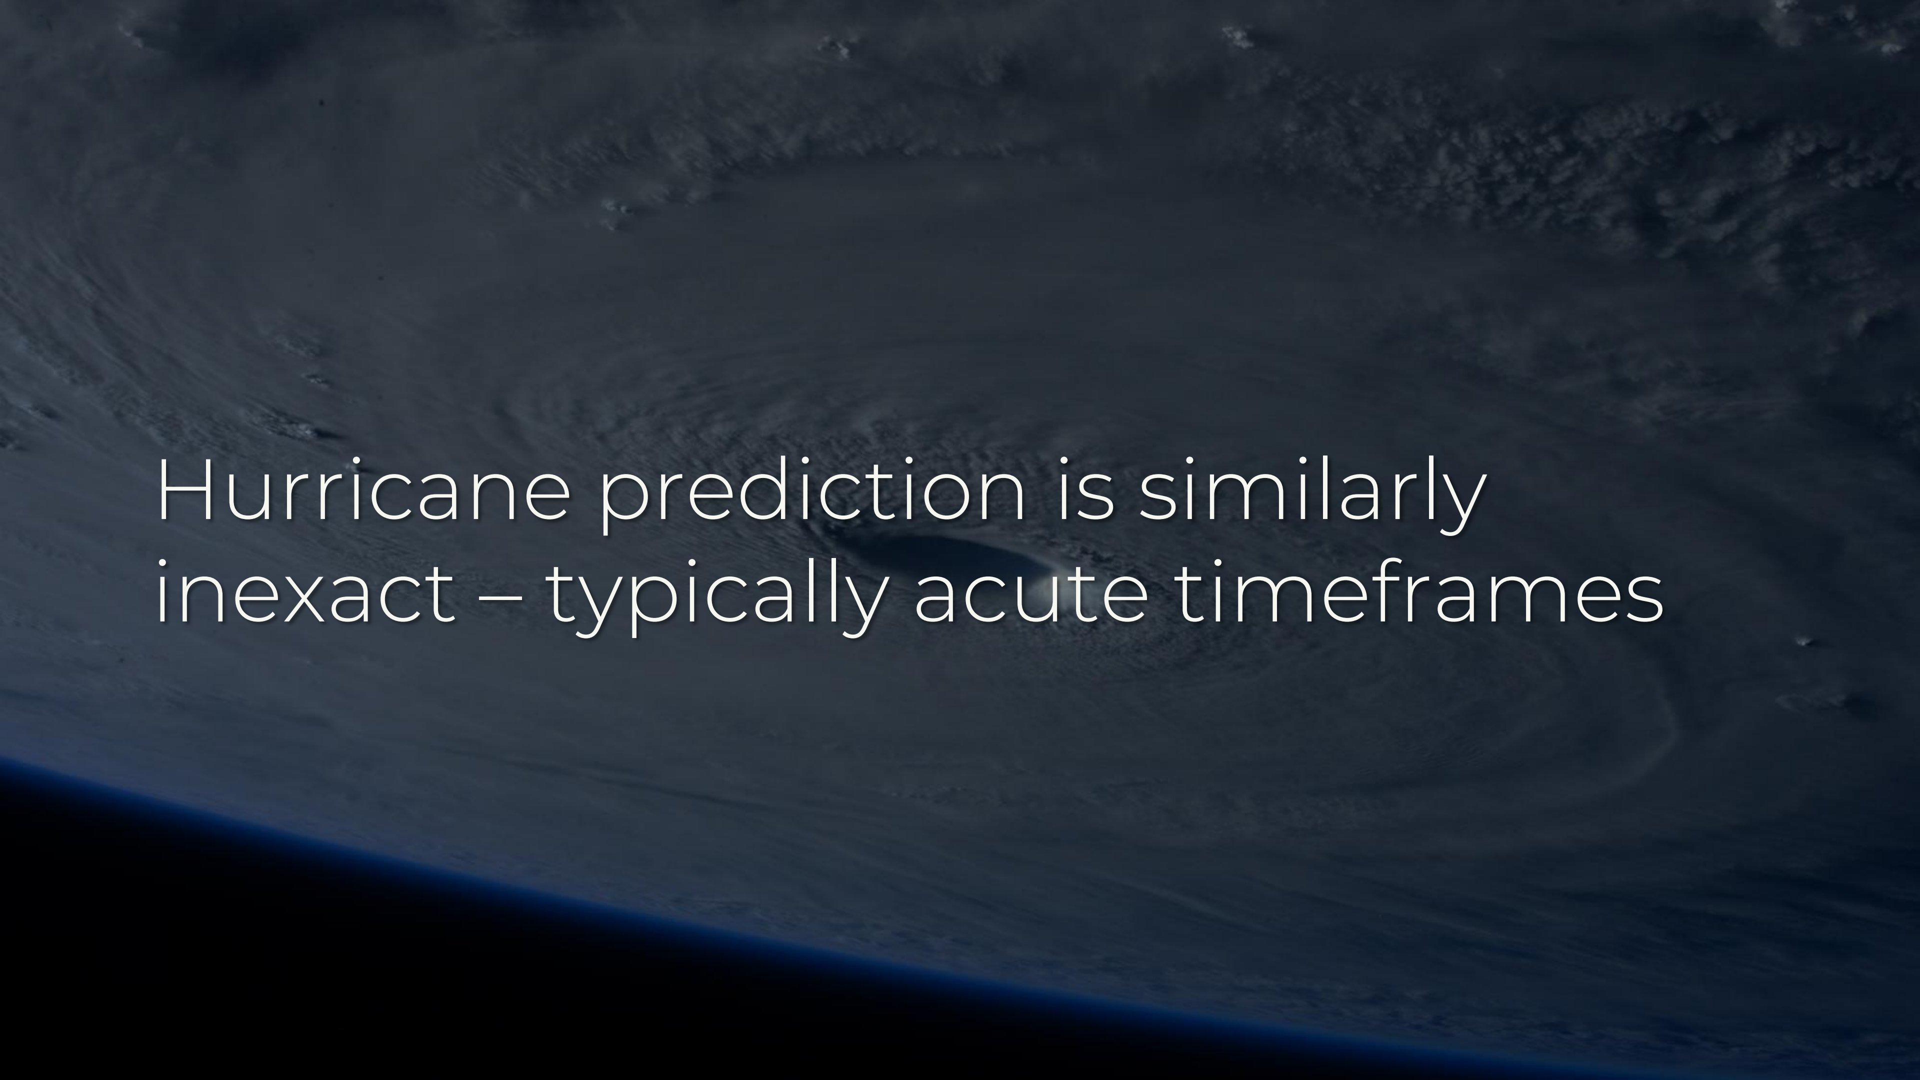

Hurricane prediction is similarly inexact – typically acute timeframes e.g. Hurricane Irma’s exact course was incorrectly predicted only days before

But we know hurricane risk zones, &
to prepare them for hurricane season

Climate change: we don’t know the exact time & sequence of events

But, we know enough to begin preparing for the most likely risks

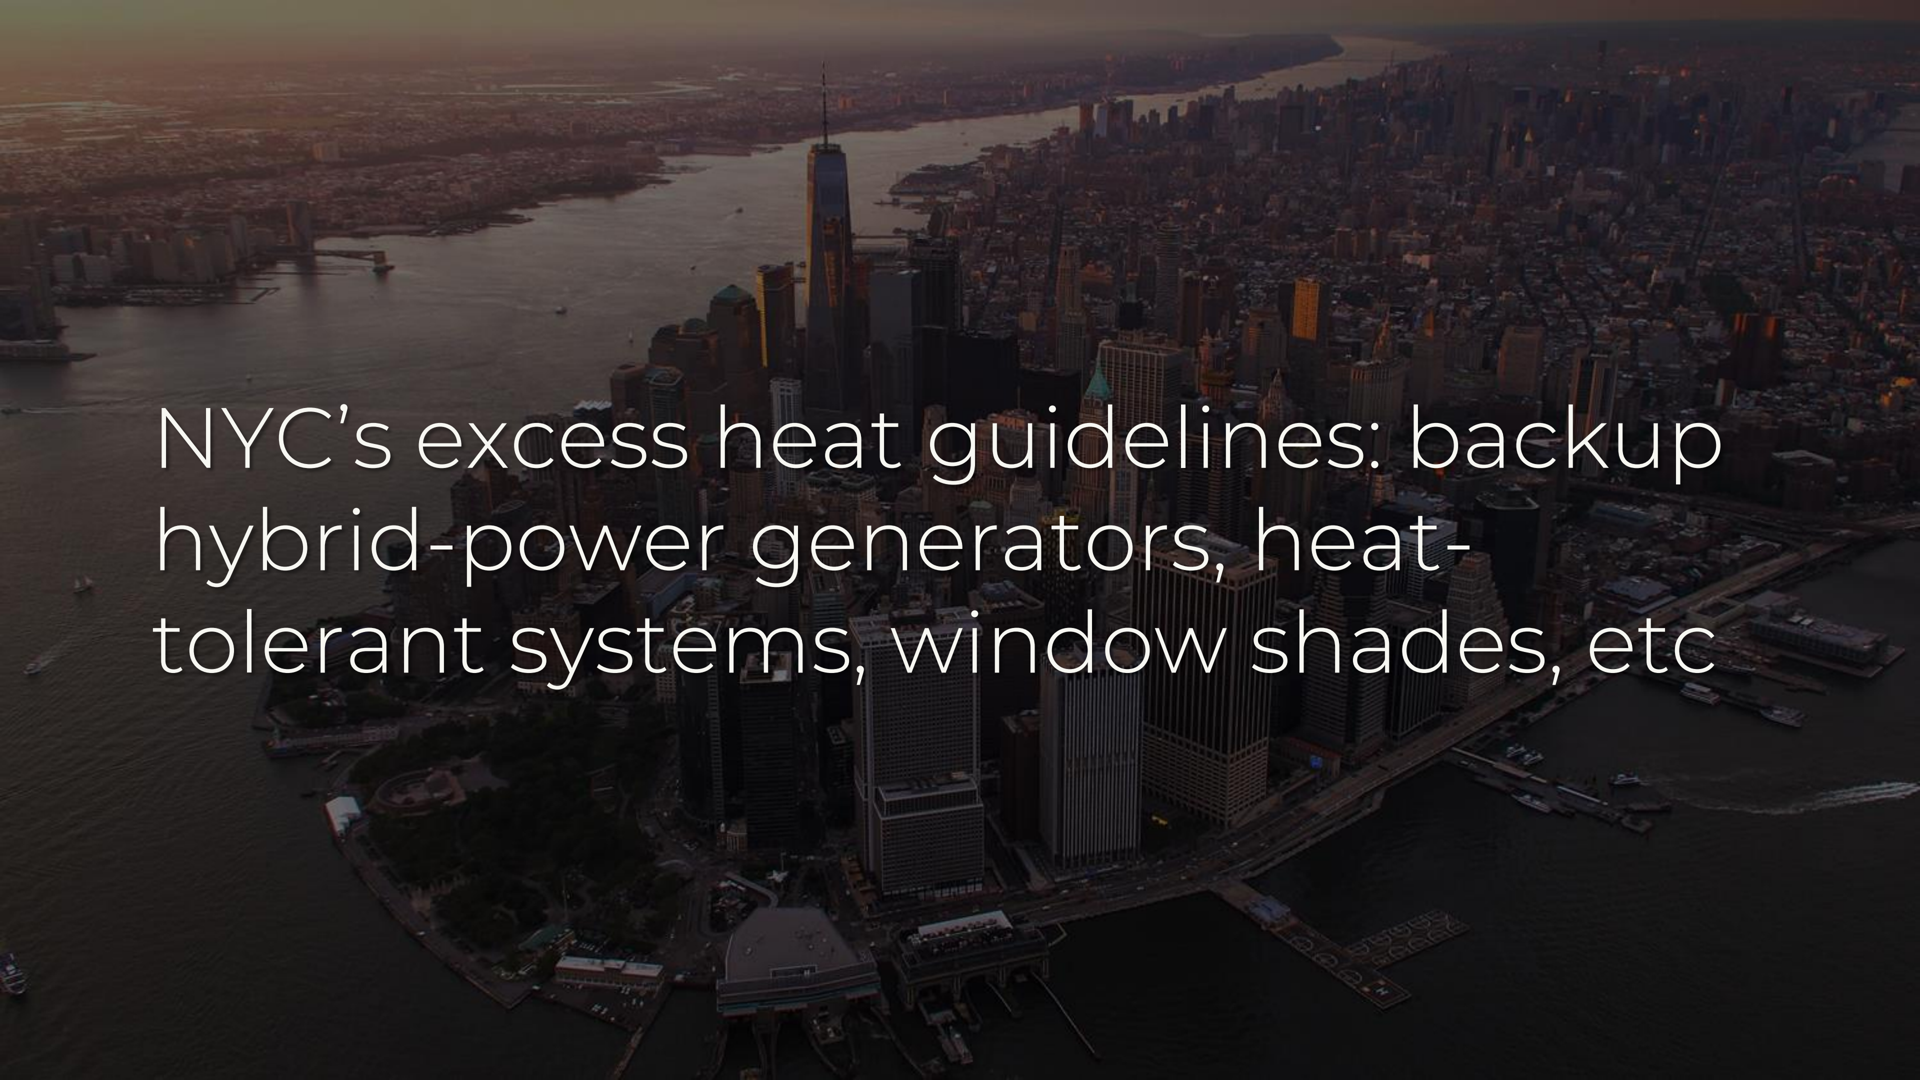

NYC’s excess heat guidelines: backup hybrid-power generators, heat-tolerant systems, window shades, etc

Financial crisis: ignoring systemic risk leads to cascading failures

Must consider common attributes that could undergo a collective shock

New Q: what is the minimum level of prediction to justify preparedness?

IMO: exact prediction is largely irrelevant – focus on hazard reduction

What should we be doing instead?

Given finite resources, it's better to research hazard assessment & reduction vs. attack prediction

WWWH&W for one attack is less valuable than knowing most probable scenario & prepping for max impact

An analogy based on a true story, via Alvaro Videla (@old_sound):

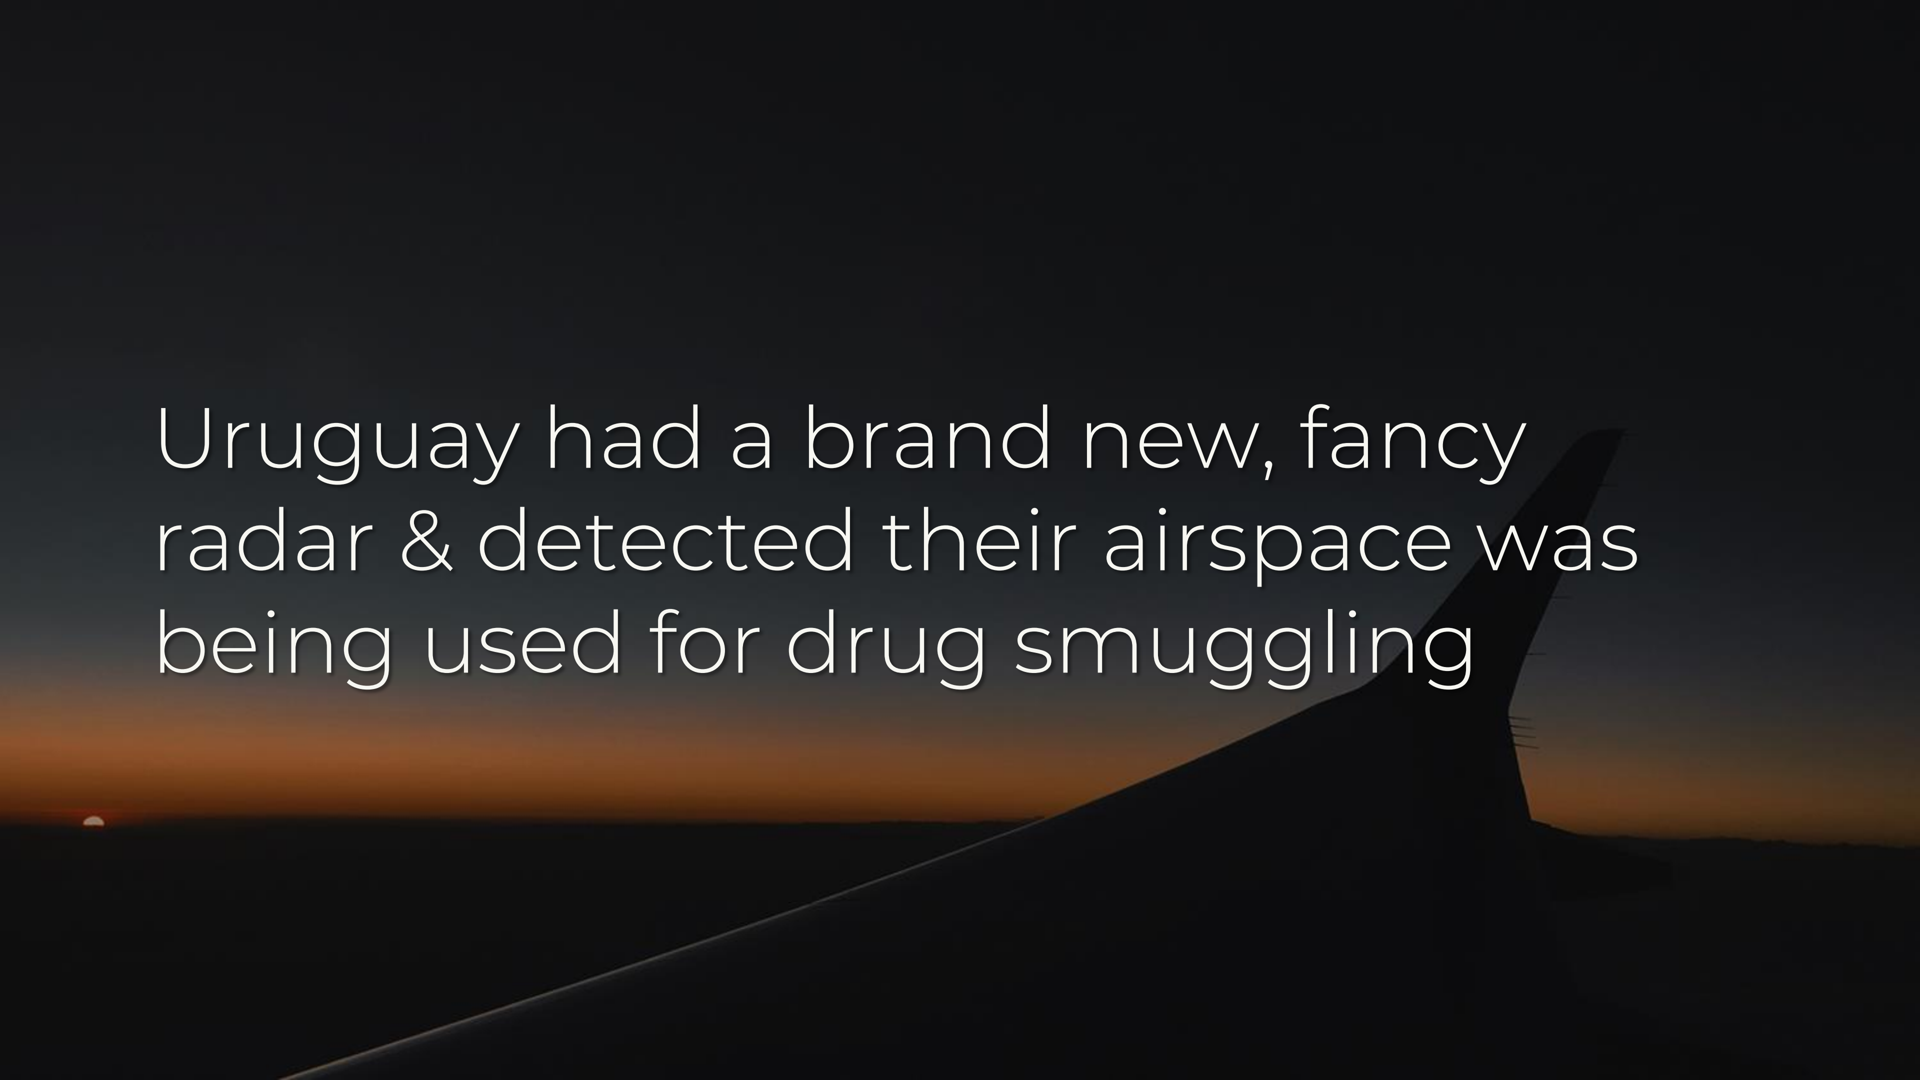

Uruguay had a brand new, fancy radar & detected their airspace was being used for drug smuggling

…but they can’t do anything because they don’t have planes fast enough to catch the bad guys

You can predict something, so what?
Can you do anything about it?

Conduct attack **forecasting** to determine general, probabilistic risk

Minimize potential impact based on business context, not security context

Step 1: Which threats actually impact business performance?

Talk to your finance colleagues about financial priorities – they won’t bite!

Anything that doesn’t disrupt revenue directly or erode “differentiation” probably doesn’t matter

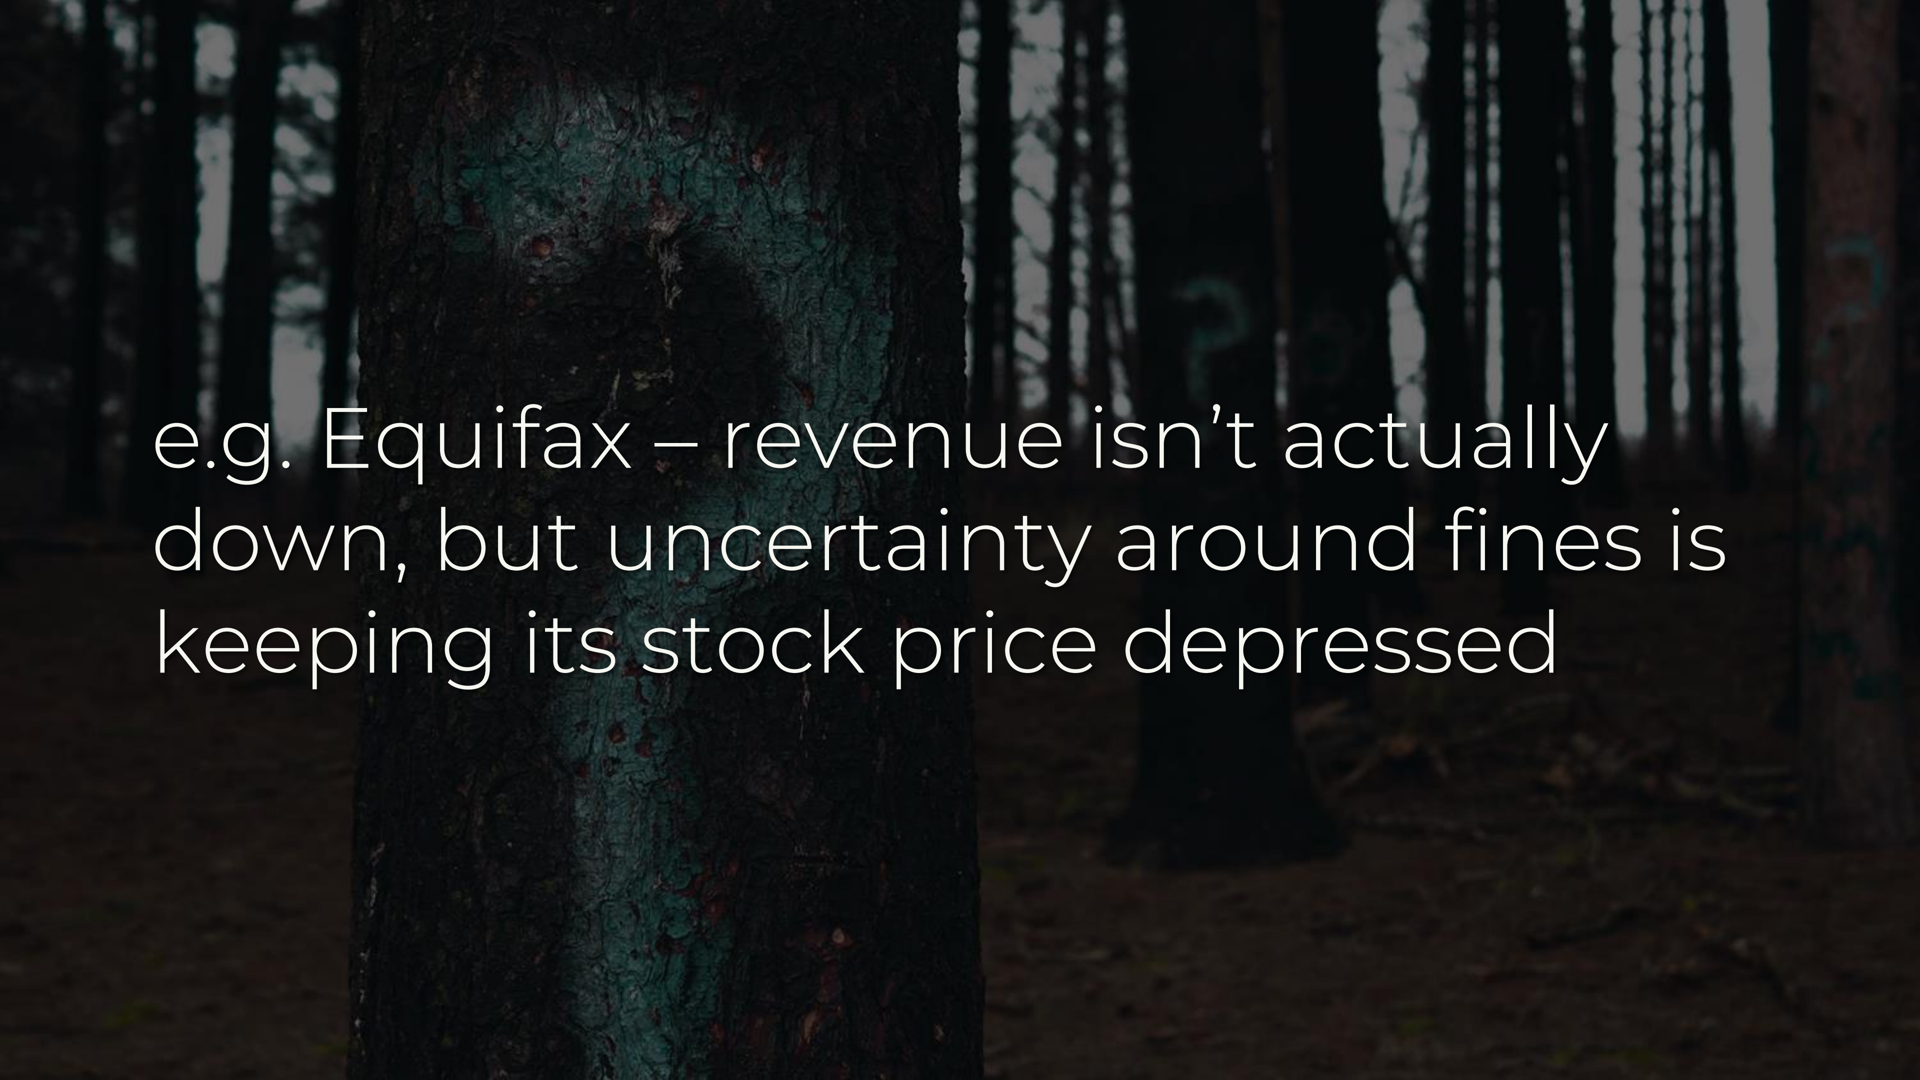

e.g. Equifax – revenue isn’t actually down, but uncertainty around fines is keeping its stock price depressed

Step 2: Assume they’ll actually happen – how can you reduce the impact?

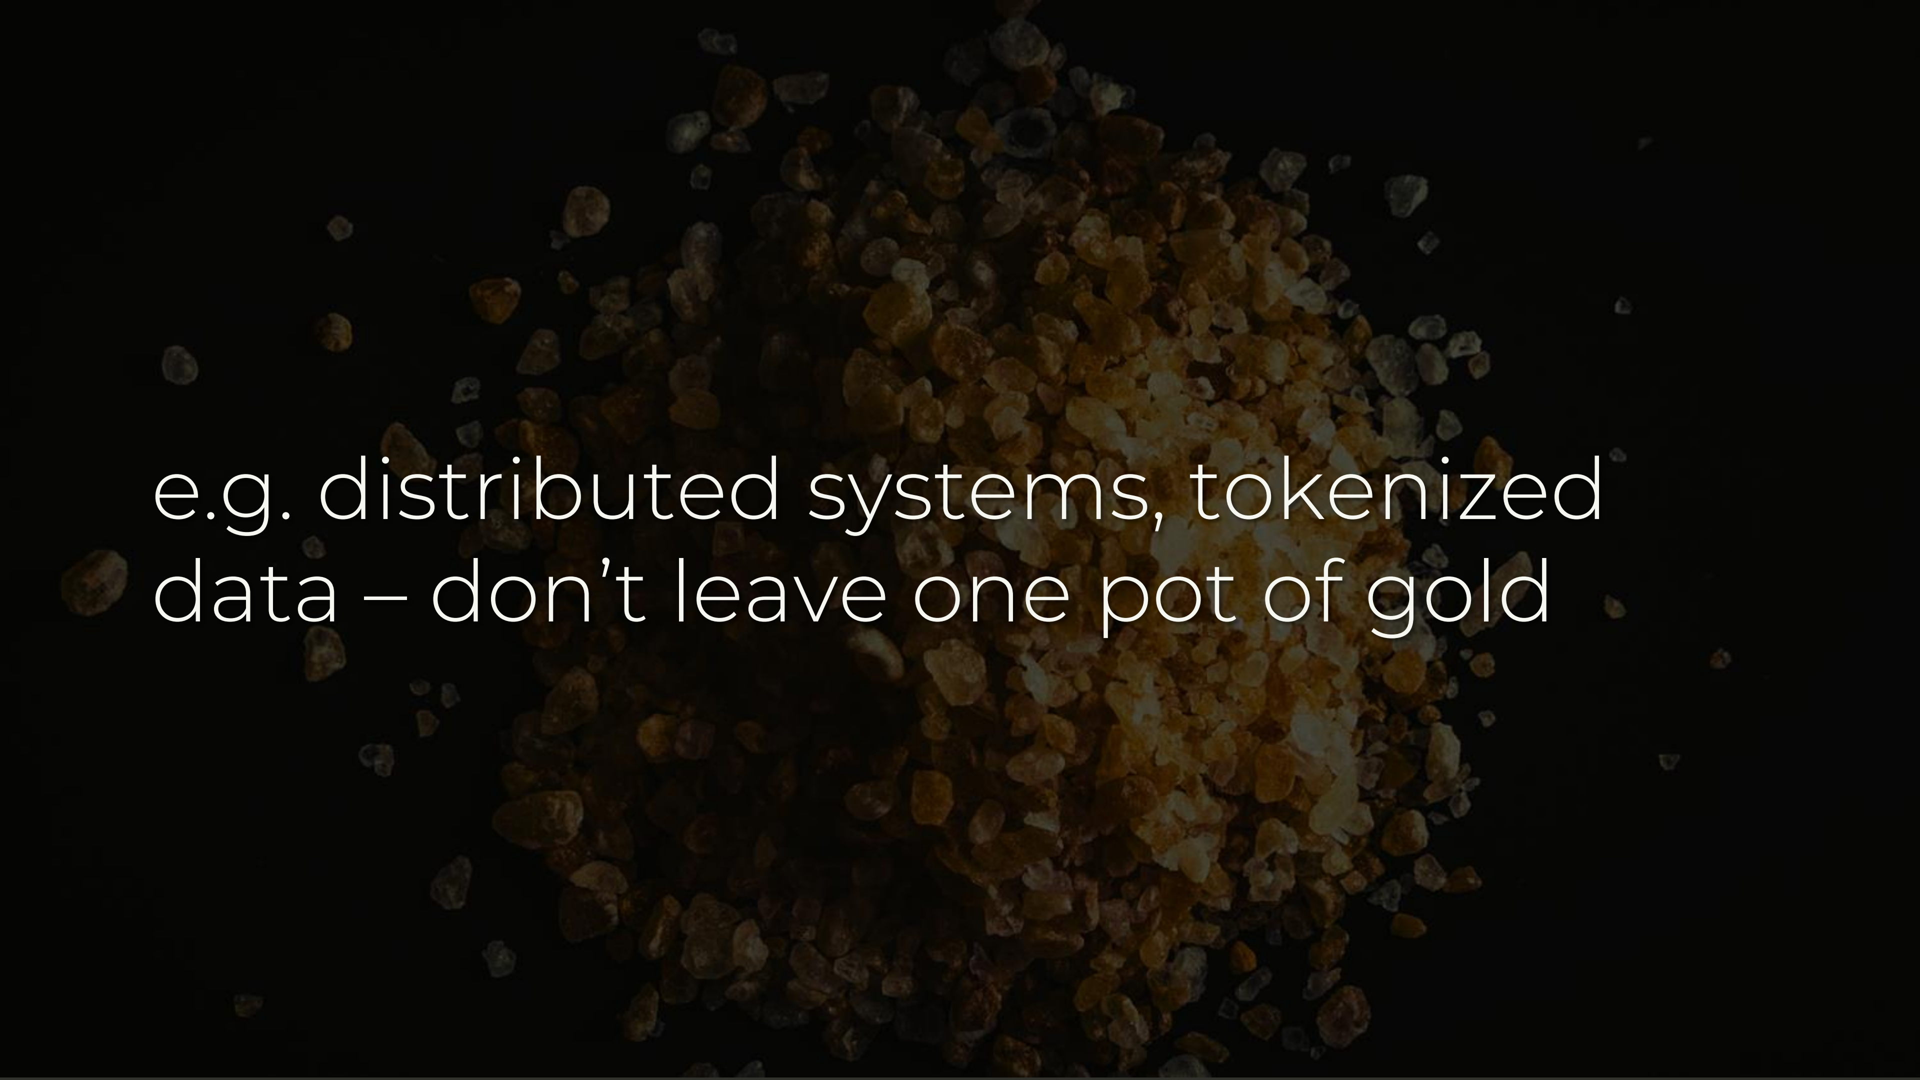

e.g. distributed systems, tokenized data – don’t leave one pot of gold

Step 3: What is an acceptable level of impact your org can tolerate?

What is material to your org? e.g. 10 mins of downtime? 60? 1440?

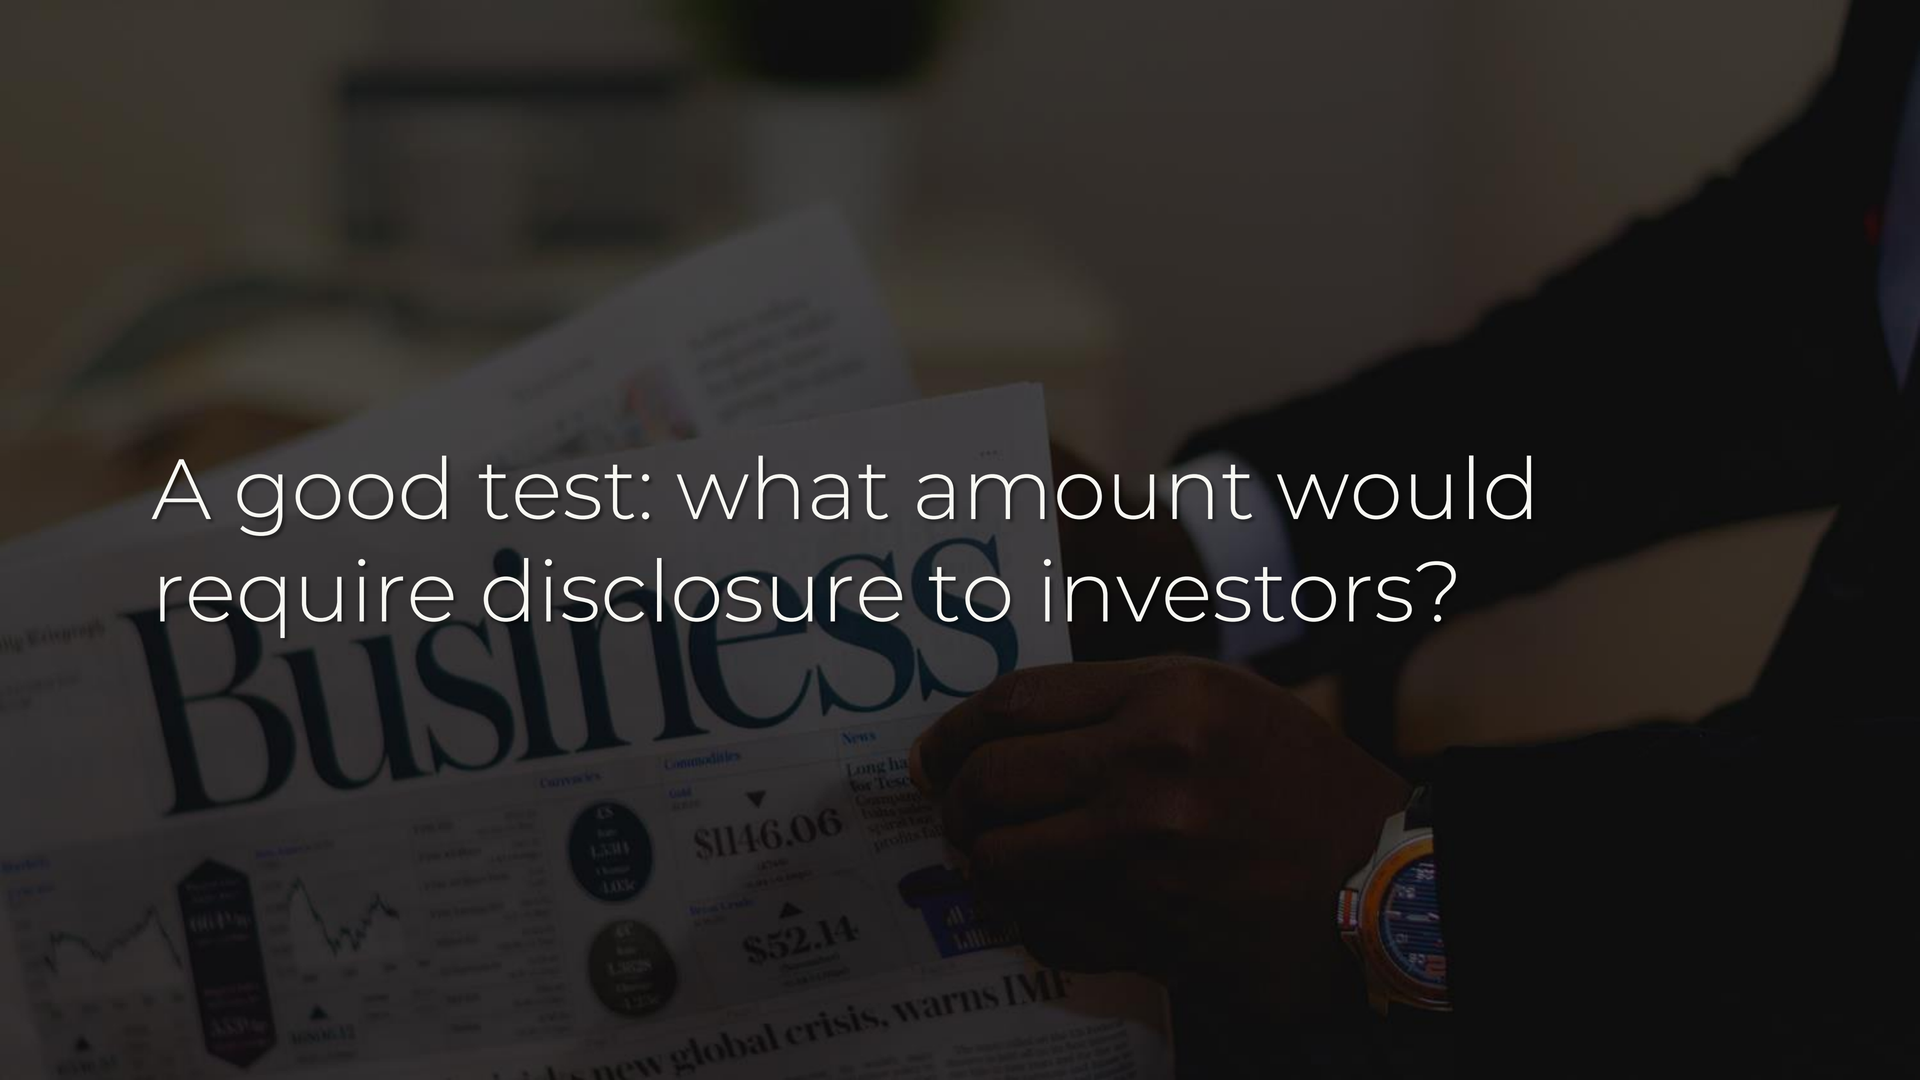

A good test: what amount would require disclosure to investors?

Exercise: How do impacts translate in $ terms? (fines, IR costs, lost revenue)

You can’t protect everything – accept some things just aren’t as important

Security teams can burn out others & themselves with “everything = critical”

e.g. Critical infrastructure: customer $ data is less important than uptime

"Resilience in infosec is a flexible system that can absorb an attack and reorganize around the threat"

– my attempt at a definition

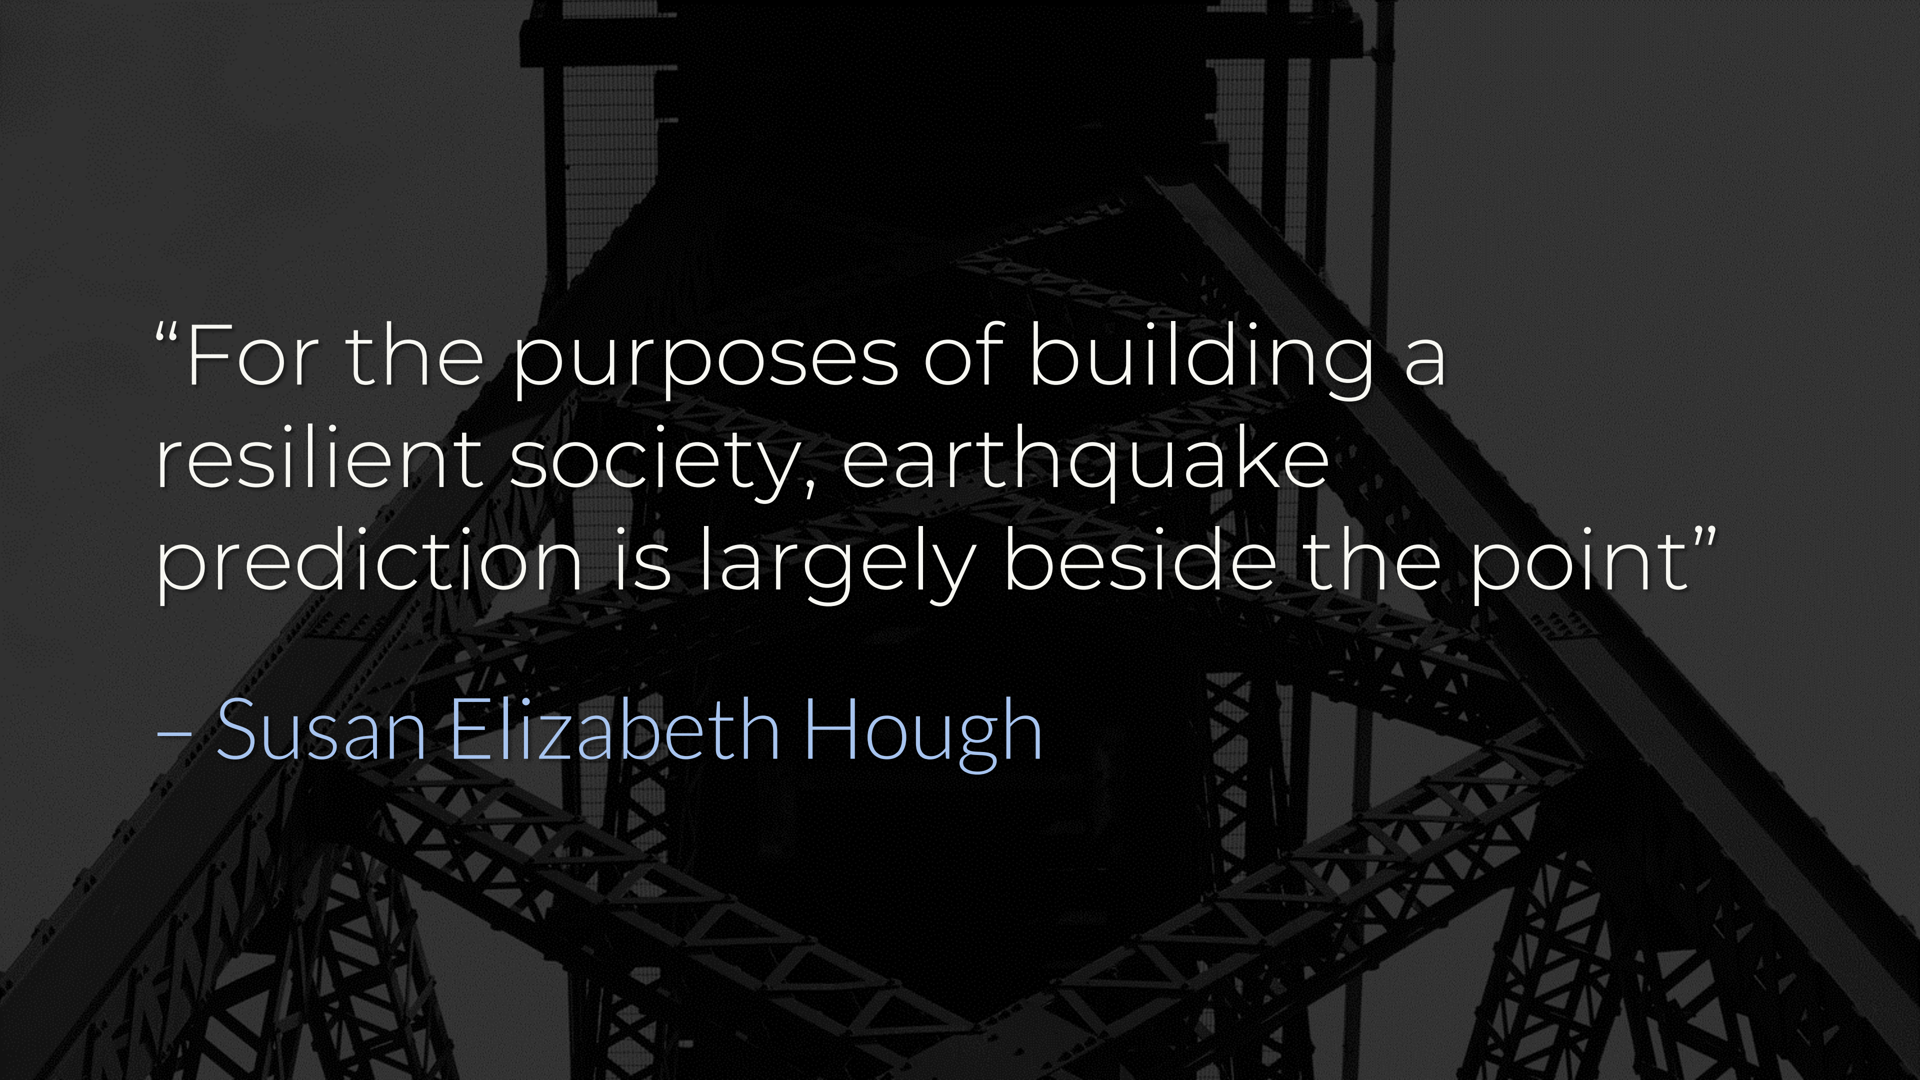

“For the purposes of building a resilient society, earthquake prediction is largely beside the point”

– Susan Elizabeth Hough

Resilience “radically accepts” an outcome & aims to reduce the hazard

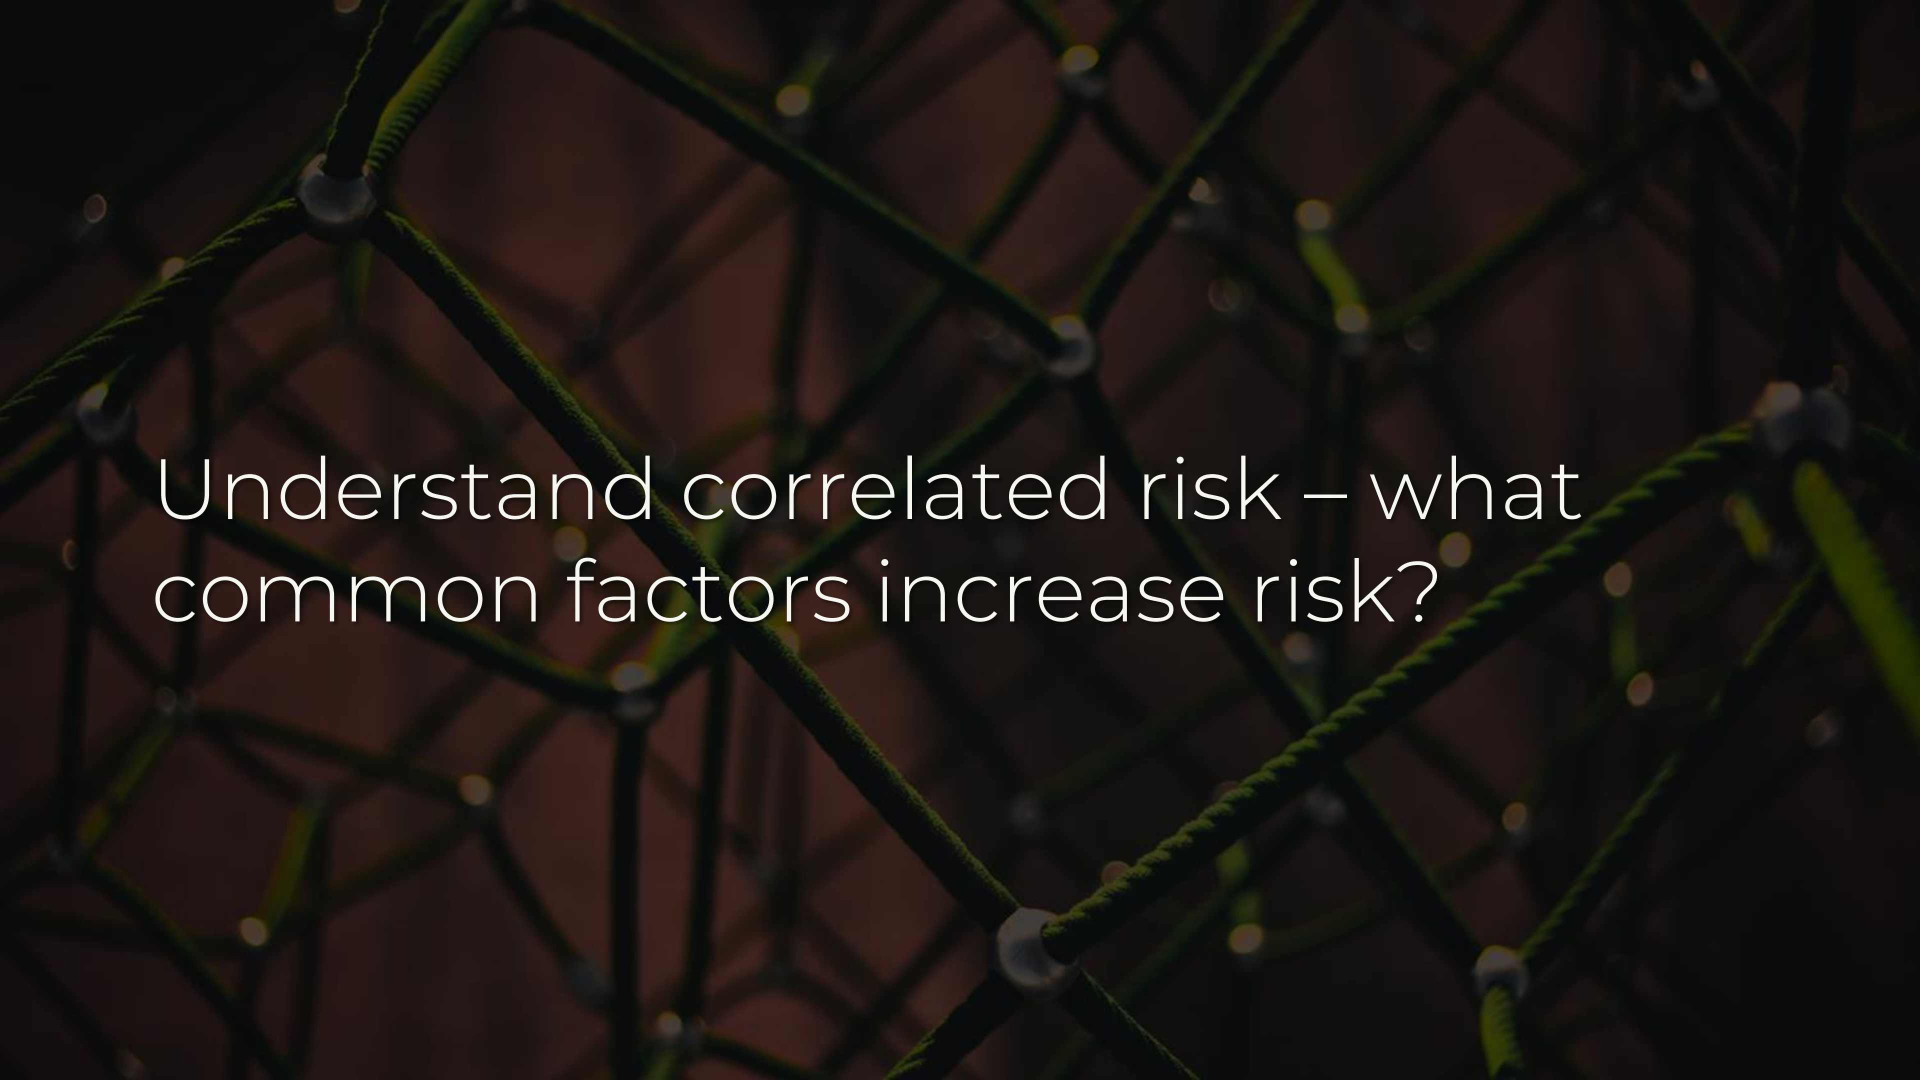

Understand correlated risk – what common factors increase risk?

Design (biz critical) systems with the assumption of compromise in mind

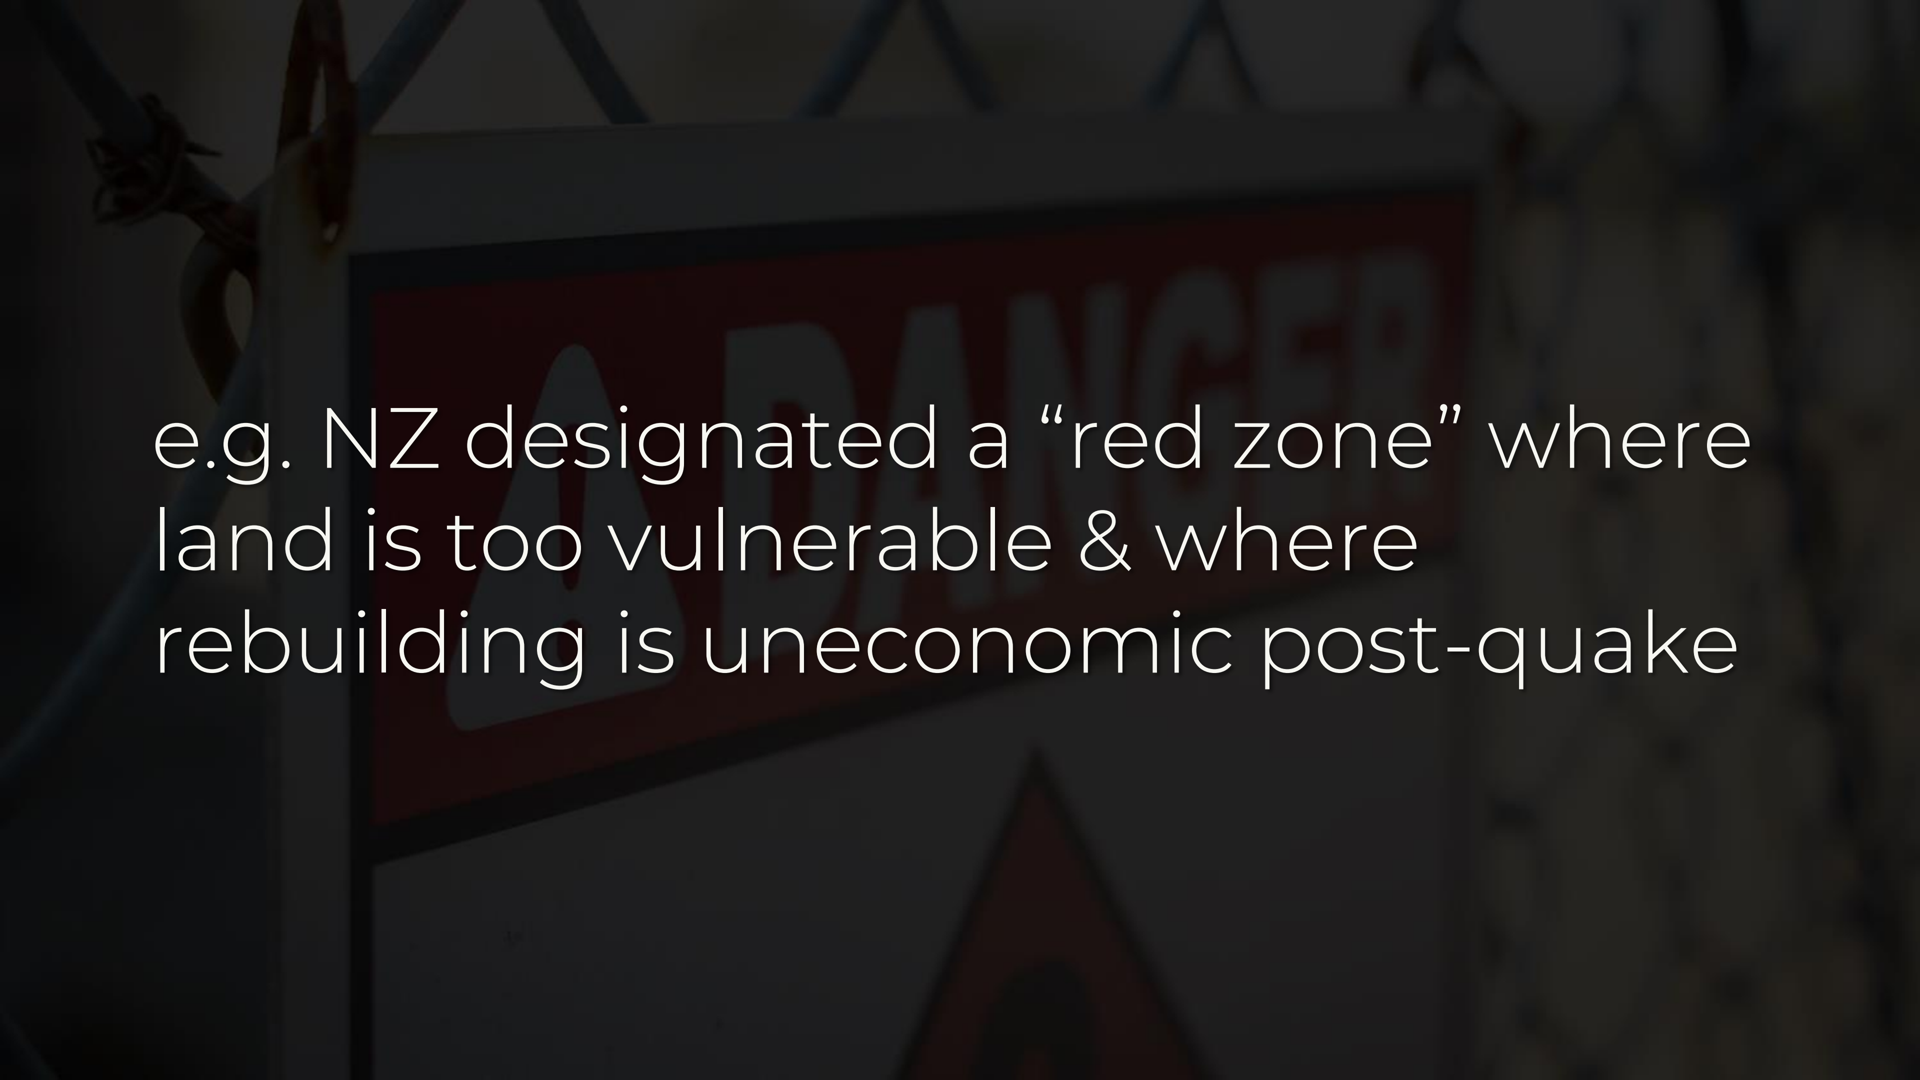

e.g. NZ designated a “red zone” where land is too vulnerable & where rebuilding is uneconomic post-quake

Identify the red zones within your IT systems (read this talk for more)

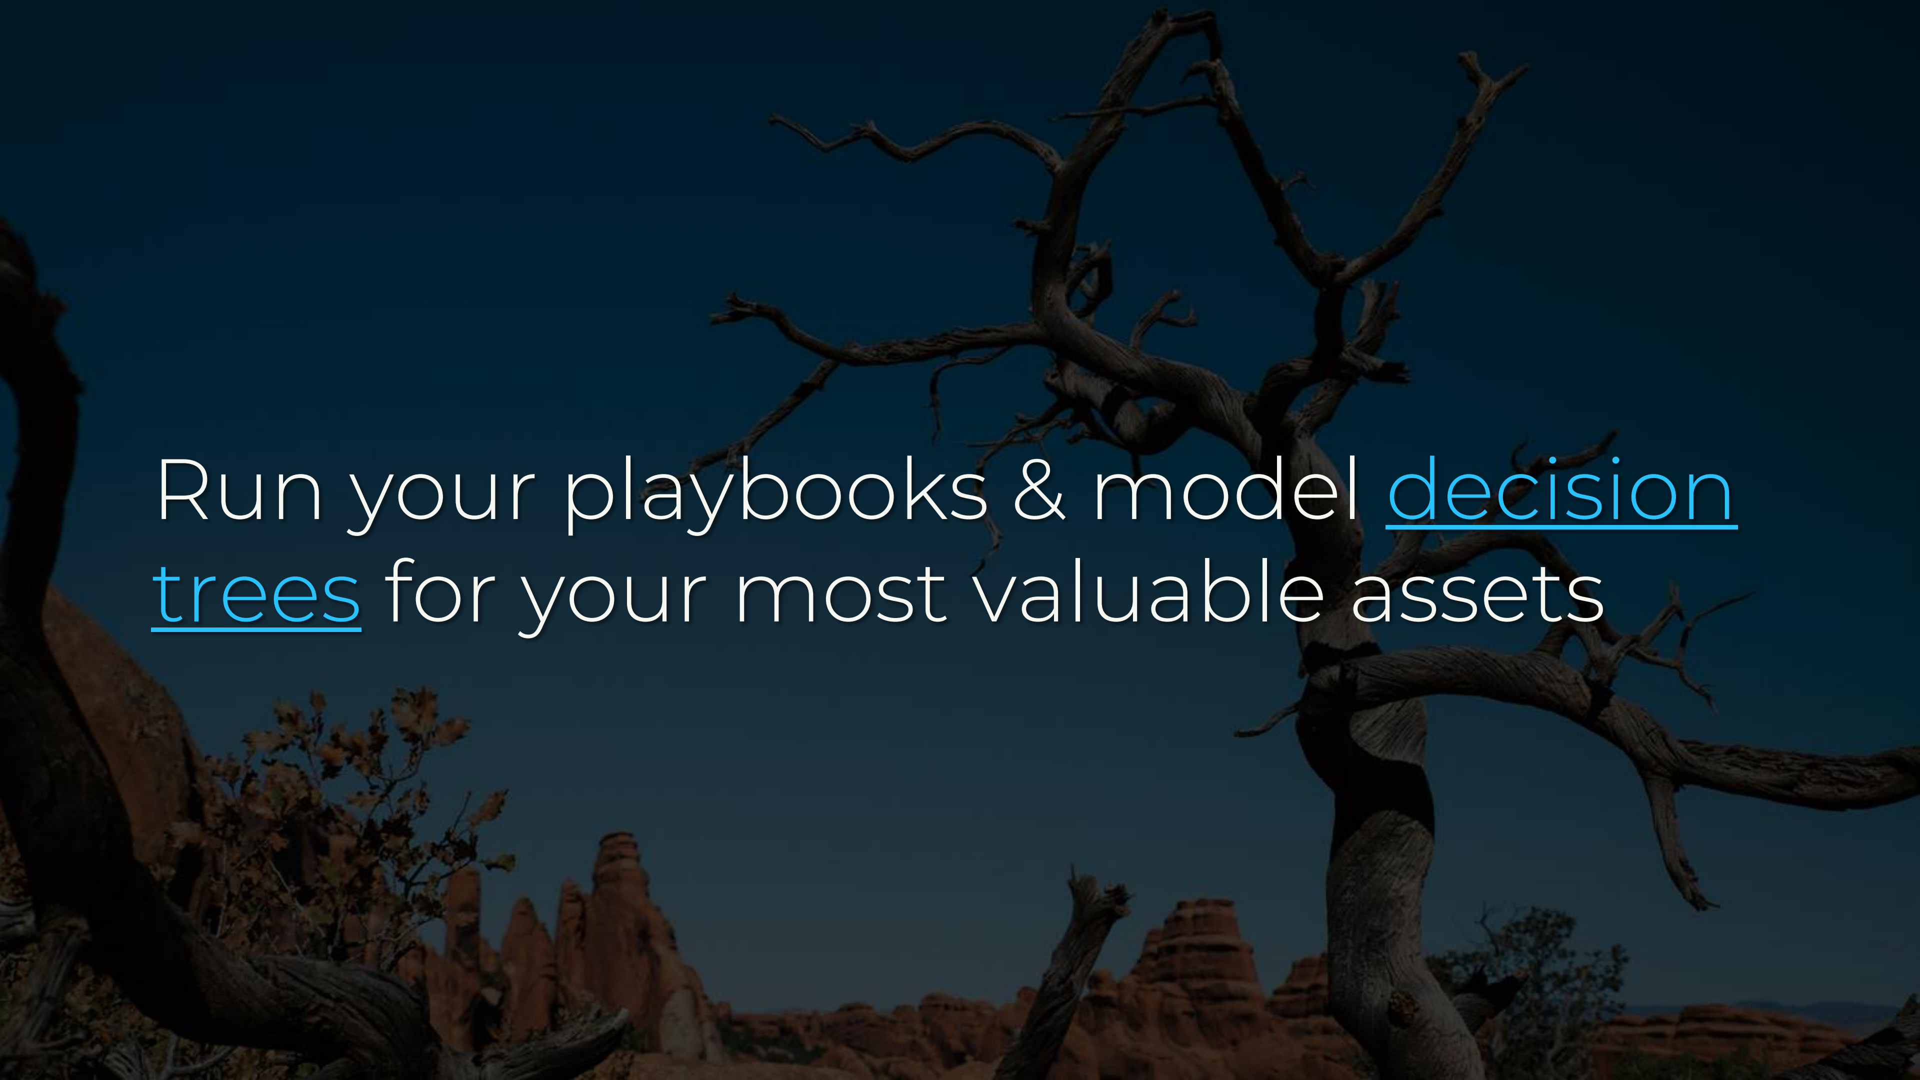

Run your playbooks & model decision trees for your most valuable assets

No point predicting if you haven’t practiced how to defend against it

Conclusion

Predicting who, when, where, how, why about an attack is unrealistic

Prediction about an individual attack is not that useful (on a relative basis)

Many “attack prediction” methods are really about detection & too myopic

Requires an inherently reactive approach – even more “things to do”

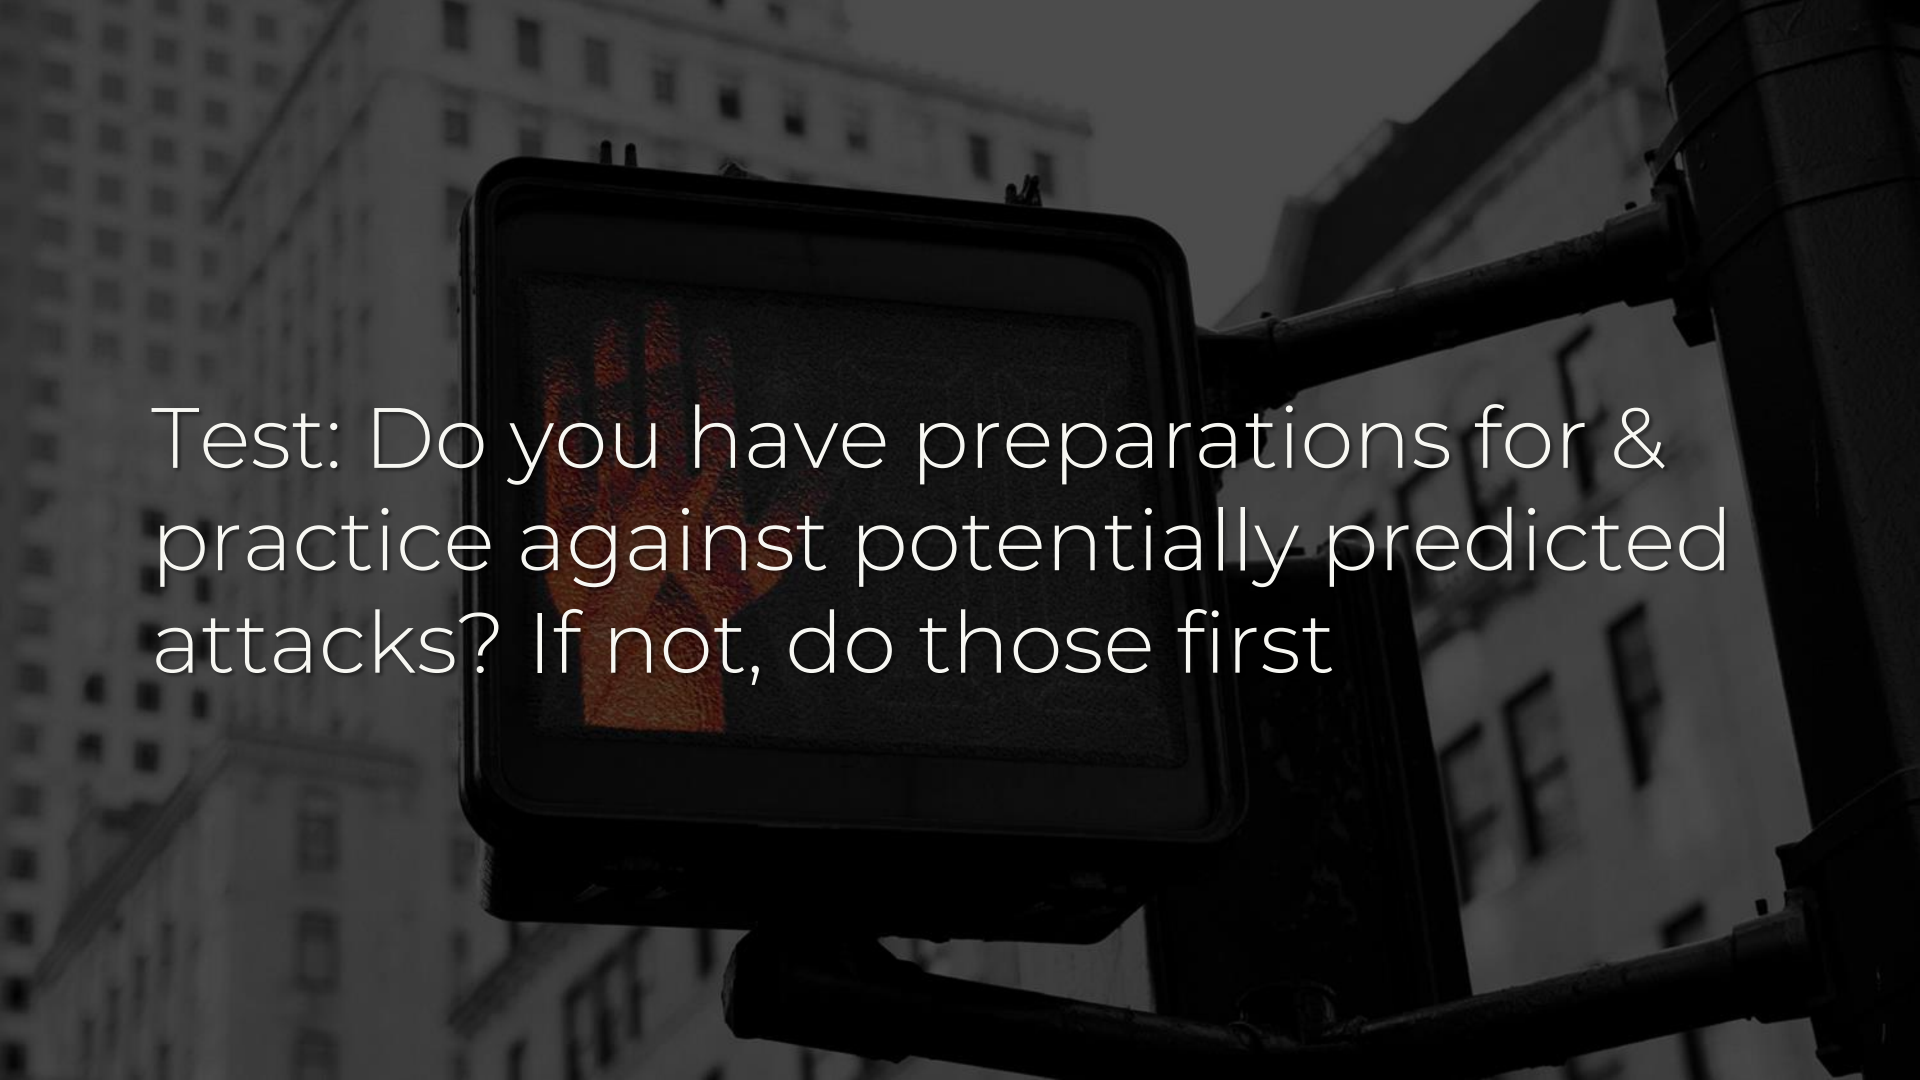

Test: Do you have preparations for & practice against potentially predicted attacks? If not, do those first

Assume pwnage & architect robust, adaptable, & transformable systems

Resilient systems support the business against many eventualities

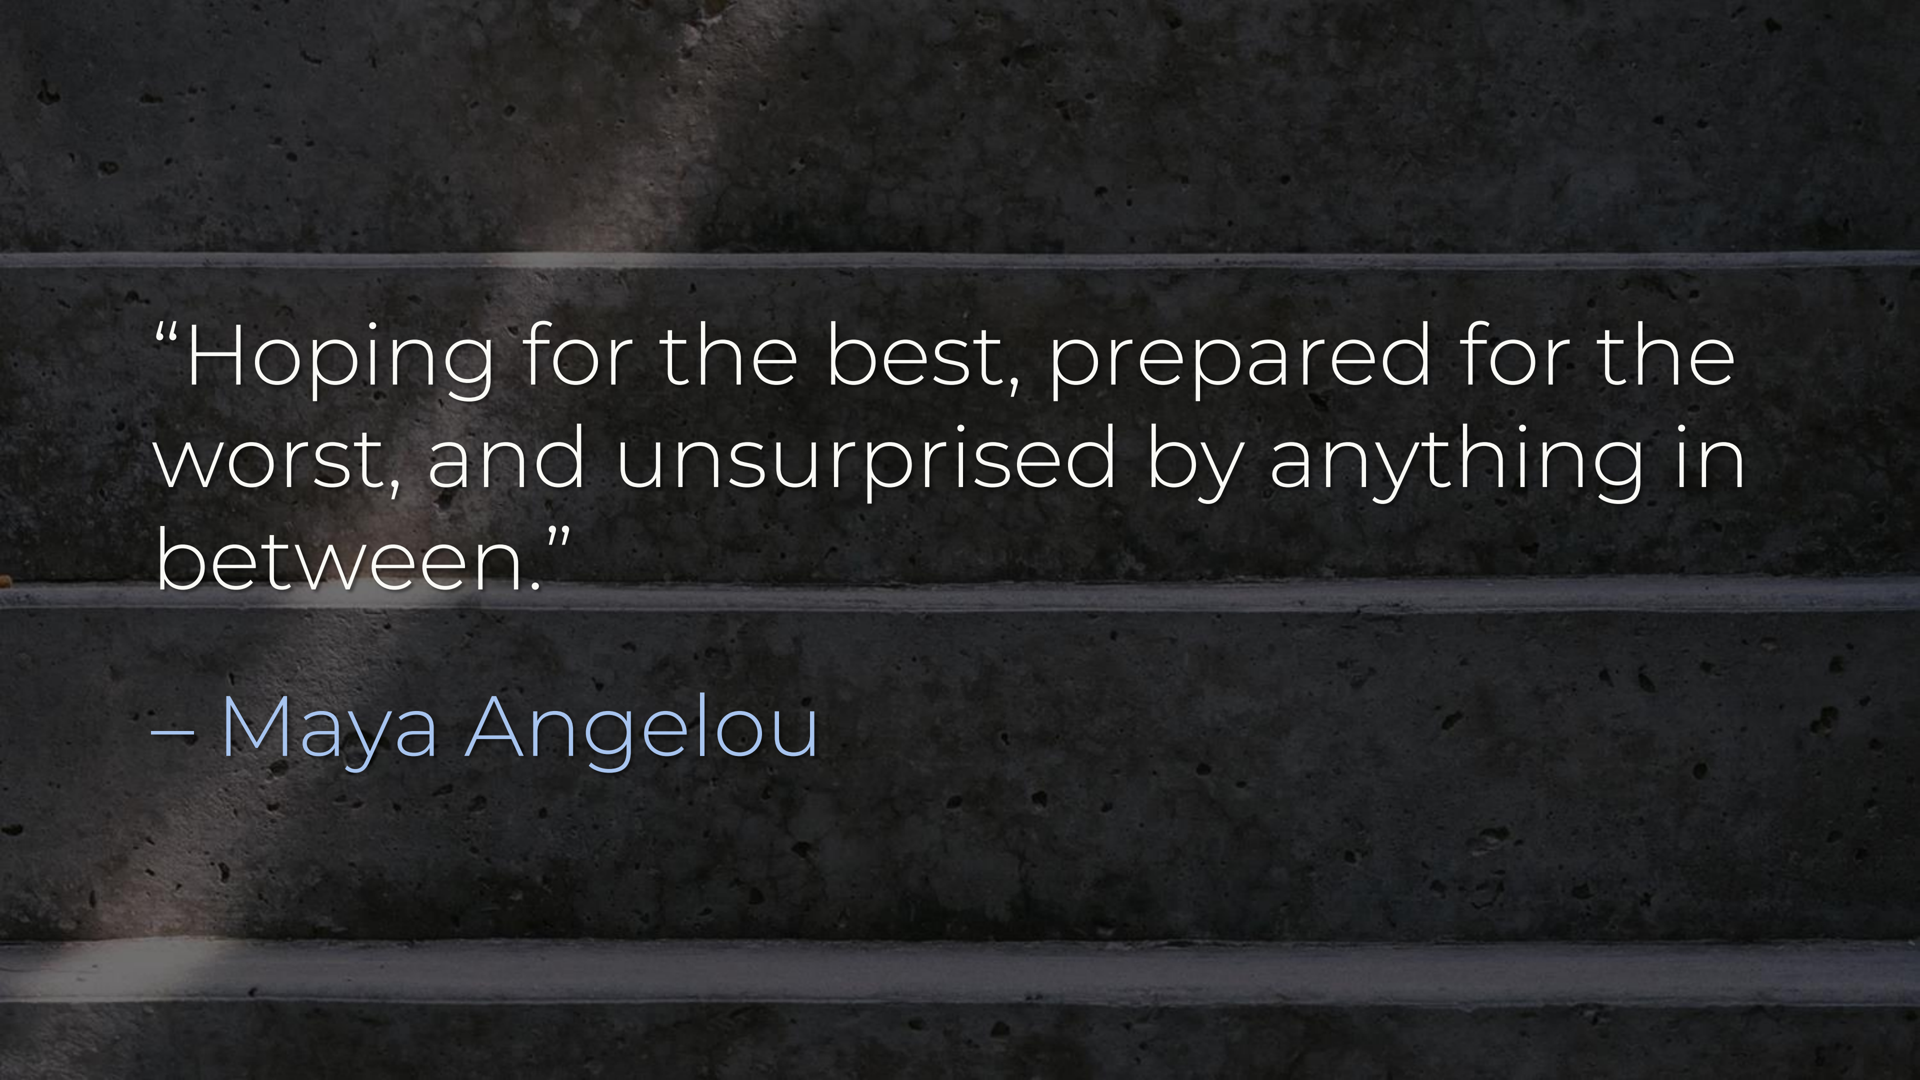

> “Hoping for the best, prepared for the worst, and unsurprised by anything in between.”
>
> – Maya Angelou

@swagitda_

/in/kellyshortridge

kelly@greywire.net