Know Thyself

Optimizing team decision-making

Kelly Shortridge

Art into Science 2017

Hi, I'm Kelly

- Now: Product Manager for Analytics at BAE Applied Intelligence
- Previously: Co-founder of IperLane; M&A banker covering infosec
- I want to make defense sexy again

What will I cover?

- Cognitive biases & their manifestations
- Group dynamics & biases
- Strategies to counter these biases
- An “easy” 6 step bias-resilience plan

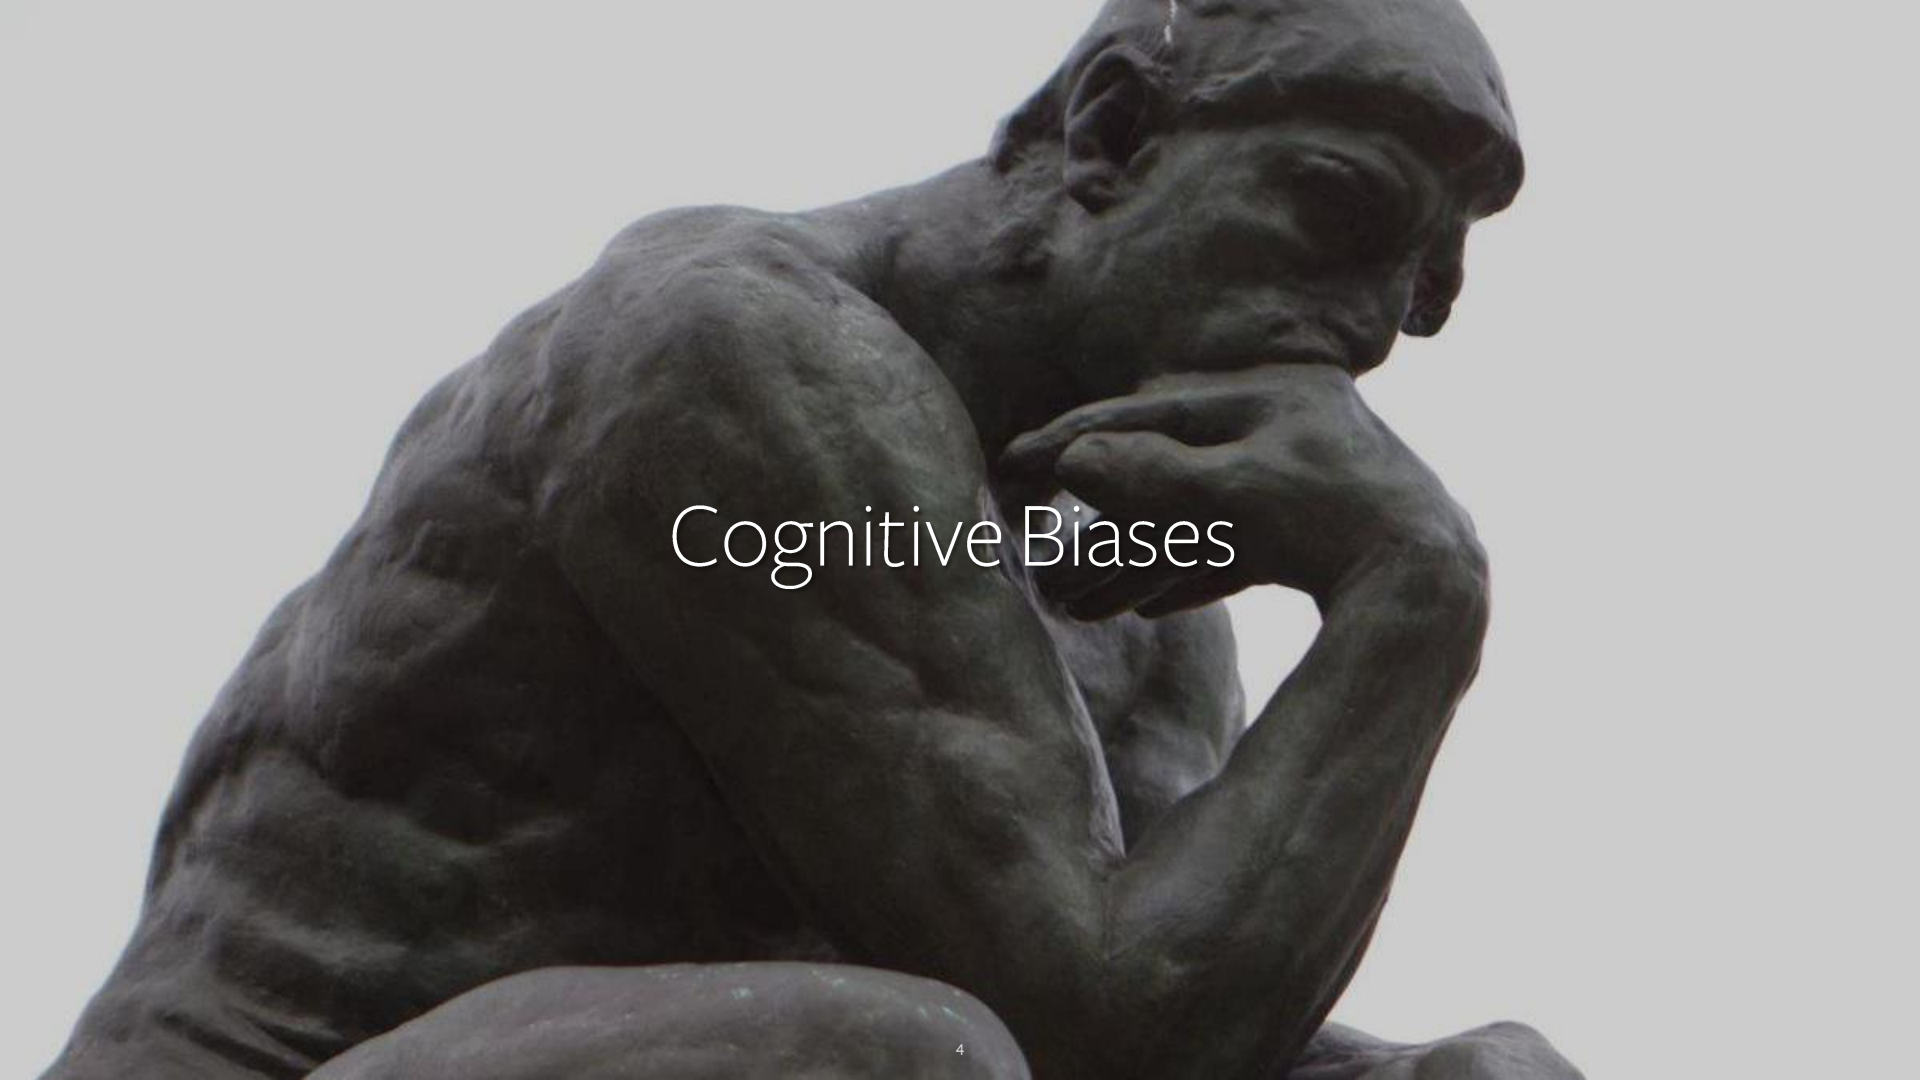

Cognitive Biases

Cognitive bias?

- Ideal = rational brain accurately weighs all potential variables and outcomes when making a decision
- In reality = “irrational” brain is fine-tuned by evolution to make speedy decisions that will help you survive
- We do not objectively evaluate input
- We create our own subjective realities

Prospect theory

- People choose by evaluating potential gains and losses via probability
- Care about relative outcomes instead of objective ones (reference point)
- Prefer a smaller, more certain gain but riskier chance of a smaller loss
- Losses hurt 2.25x more than gains feel good
- Overweight small probabilities and underweight big ones
- Diminishing sensitivity to losses or gains the further away from ref point

Offense vs. Defense

Offense

- Risk averse
- Quickly updates reference point
- Focus on probabilistic vs. absolute outcome

Defense

- Risk-seeking
- Slow to update reference point
- Focus on absolute vs. probabilistic outcome

Prospect theory in InfoSec

- Defenders overweight small probability attacks (APT) and underweight common ones (phishing)
- Defenders also prefer a slim chance of a smaller loss or getting a “gain” (stopping a hard attack)
- Attackers avoid hard targets and prefer repeatable / repackagable attacks (e.g. malicious macros vs. bypassing EMET)

Other biases in infosec

- Time inconsistency: current self = different than future self
  - We don't want to do things that have a delay to the reward, even if the reward is bigger (Marshmallow Experiment)
  - Technical debt in a nutshell & perpetuates cat & mouse game
- Dual System Theory: mind system 1 ("lizard brain") = automatic, fast, non-conscious, mind system 2 = controlled, slow, conscious
  - System 1 often dominant in decision-making, particularly with pressure
  - System 1 = flashy demos & sexy word salads, known strategies & products, cares about ego & succumbs to fear

What are the outcomes?

- Criminally under-adopted (corporate) tools: EMET, 2FA, canaries, white-listing, thinking along the entire killchain
- Criminally over-adopted tools: prevention tools, delivery-stage-only IDS, uncontextualized threat intel, dark-web anything
- Like having lots of firefighters, a concrete door with a heat sensor & lots of info on how fires can be started… but inside you have wooden furniture, open windows and no smoke alarms

An outcome

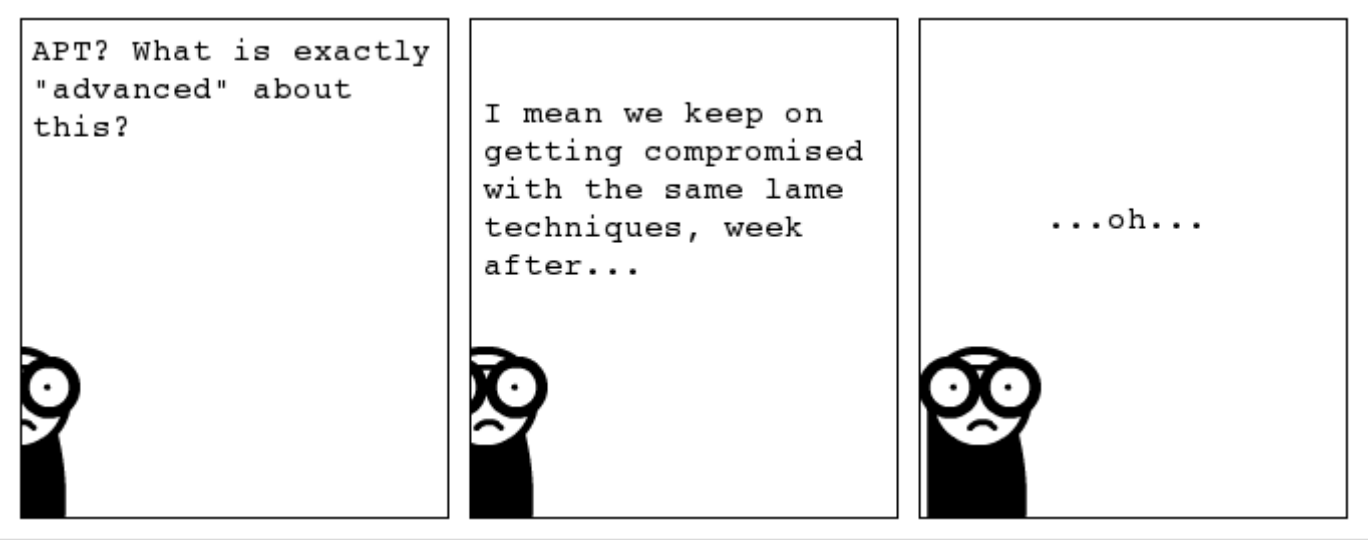

Incentive problems

- Defenders can't easily evaluate their current security posture, risk level, probabilities and impacts of attack
- Defenders only feel pain in the massive breach instance, otherwise "meh"
- Attackers mostly can calculate their position; their weakness is they feel losses 3x as much as defenders
- The high-stakes nature of the job facilitates System 1 thinking


Group Dynamics

Cognitive biases in groups

- A leader creates new social issues – if the leader's biases are stated before a discussion, that tends to set the decision
- Some evidence that groups have a stronger "escalation of commitment" effect (doubling down)
- The term "groupthink" exists for a reason
- Groups are potentially even better at self-justification, as each individual feels the outcome is beyond their control

Bosses & risk

- Boss = awareness that skill level is being evaluated
- Risky decisions make subordinates appear more competent
  - Expectation of failure = look better if it succeeds and have no penalty if it doesn't
- Fear of appearing incompetent
  - Expectation of success = penalty if it fails, not much benefit if succeeds

Example security org split

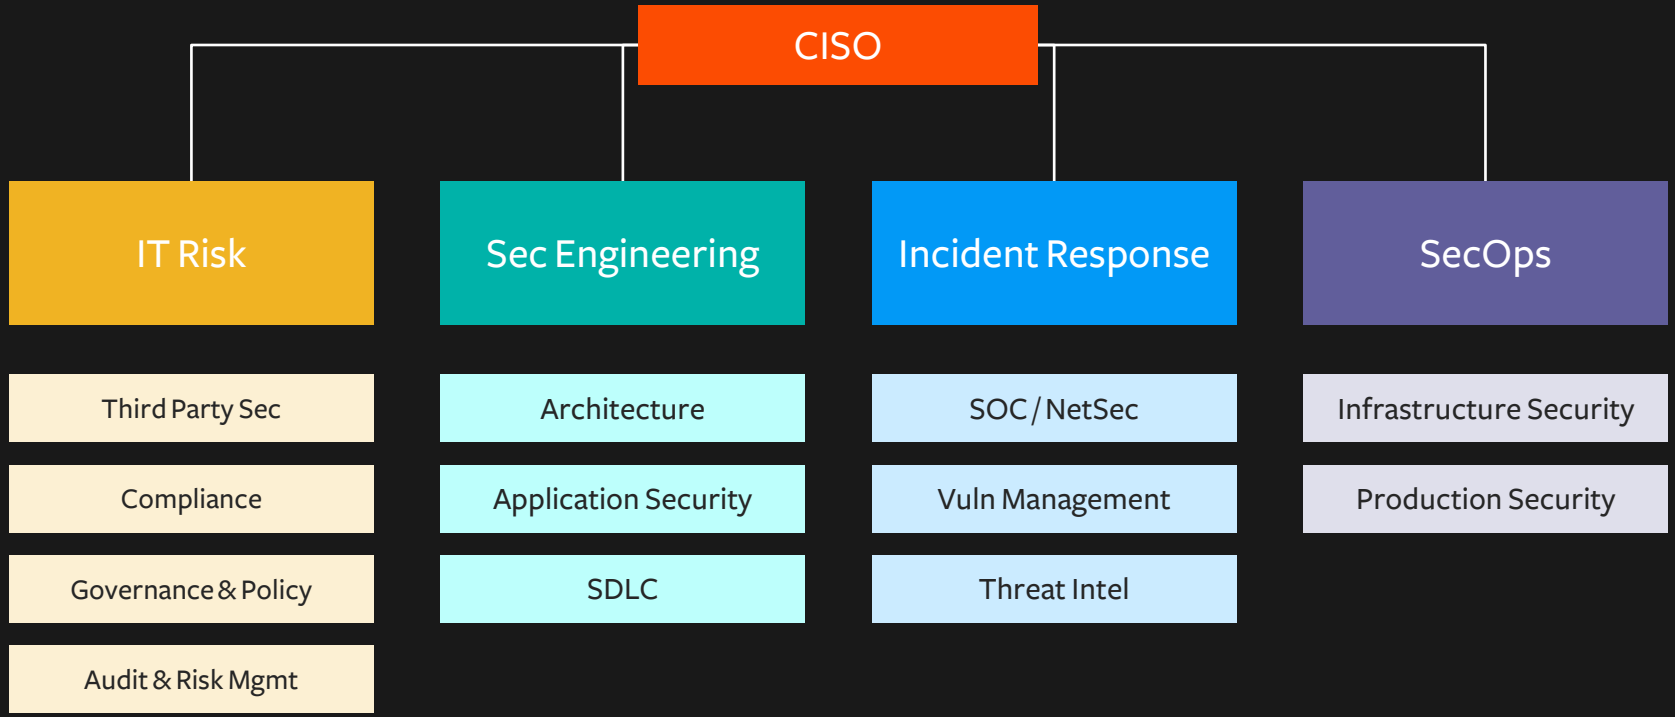

The setting

- CISO overlooks managers, who overlook an often relatively flat team
- Everyone in the security organization wants CYA – they're first in line to be blamed in event of a breach
- No one will ever get security 100% correct (some failure is assumed)
- Viewed as a cost-center
- Non-managers often become box-minders, regardless of role

How are CISOs evaluated?

- Reducing costs, delivering projects on time, increasing efficiency
- ...oh and also minimizing the company's risk profile
- Ability to sell the "vision" & communicate with CEO, CFO, COO, Board
- Responsible for managing any security incidents (during & after)
- Thought-leadering in the community

Success & failure for blue team members

- For their boss
  - Success = helping reduce cost, deliver on time, increase efficiency
  - Failure = a breach, increasing costs, slow delivery
- Defending against super sick APT = expectation of failure (ROI looks better)
- Defending against skiddies = expectation of success (ROI looks worse)
- Improving security often at odds with lower costs or faster delivery

Being a cost center adds to the issues

- Cost center = harsher penalty with screw-ups, less reward for success
- Also incentivizes creating “wow” moments to prove value
- Sunk cost fallacy is rampant – less room to admit something isn’t working and switch to something else
- Moonshot projects are reserved for revenue-generators – hard to argue for longer-term, lower-risk projects with delayed payoff

A sample meeting

- Boss proposes sticking with current plan
- Team member wins if they propose something to reduce costs or speed up delivery, or to make it seem sexier
- Team member loses if they disagree with the group, or propose something that takes more time, or money (at least short-term)
- Boss tells team member to do a risky thing, agrees to it so they don't seem incompetent

Current decision making process

- Putting out fires first, then risk mitigation (emergency room + first responders)
- Often reactionary vs proactive
- Ad-hoc brainstorming
- Focus on compliance
- Enumerating best practices


Strategies
(now entering the realm of decision trees)

Belief prompting & hard metrics

- Ask for explicit beliefs about what their opponents will do & who they are
  - Assumptions around their capital, time, equipment, risk aversion
- Model decision trees both for offense and defense
  - Use kill chain as guide for offense's process
- Theorize probabilities of each branch's outcome
  - Phishing is far more likely the delivery method than Stuxnet-style
  - Creates tangible metrics to deter self-justification

Example belief prompting

- “How do you think our adversary chooses their delivery method?”
- “What countermeasures do they anticipate?”
- “Which of our assets will attackers want?”
- Generally, for each move, map out:
  - (Defensive) How would attackers pre-emptively bypass the D move?
  - (Defensive) What will they do next in response to the D move?
  - (Offensive) Costs / resources required for the O move?
  - (Offensive) Probability the O move will be conducted?

A relevant thought leader quote

"Attackers will take the least cost path through an attack graph from their start node to their goal node"

– Dino Dai Zovi, "Attacker Math"

Examples of belief prompting

- Should we use anti-virus or whitelisting?
  - Adds recon step of figuring out which apps are on whitelist
  - Requires modifying malware so it isn't caught by an AV signature
  - Latter is way easier / cost-effective, so more likely to use it

- Skiddie randomly lands on one of our servers, what do they do next?
  - Perform local recon, escalate to whatever privs they can get
  - Counter: priv separation, don't hardcode creds
  - Leads to: attacker must exploit server, risk = server crashes

Example progression: Exfiltration

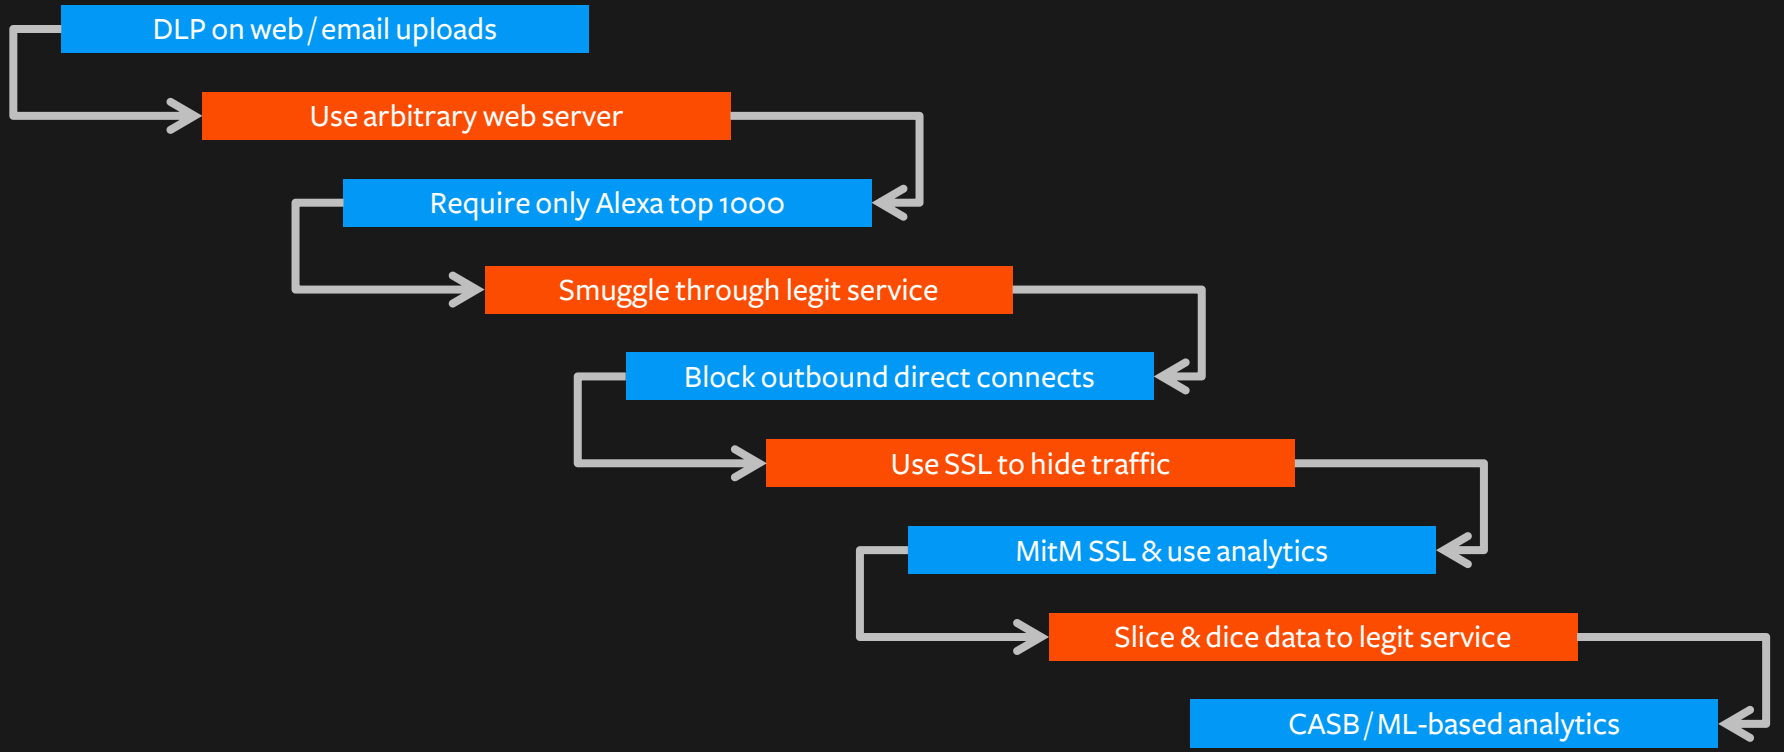

Example AD tree (for illustrative purposes)

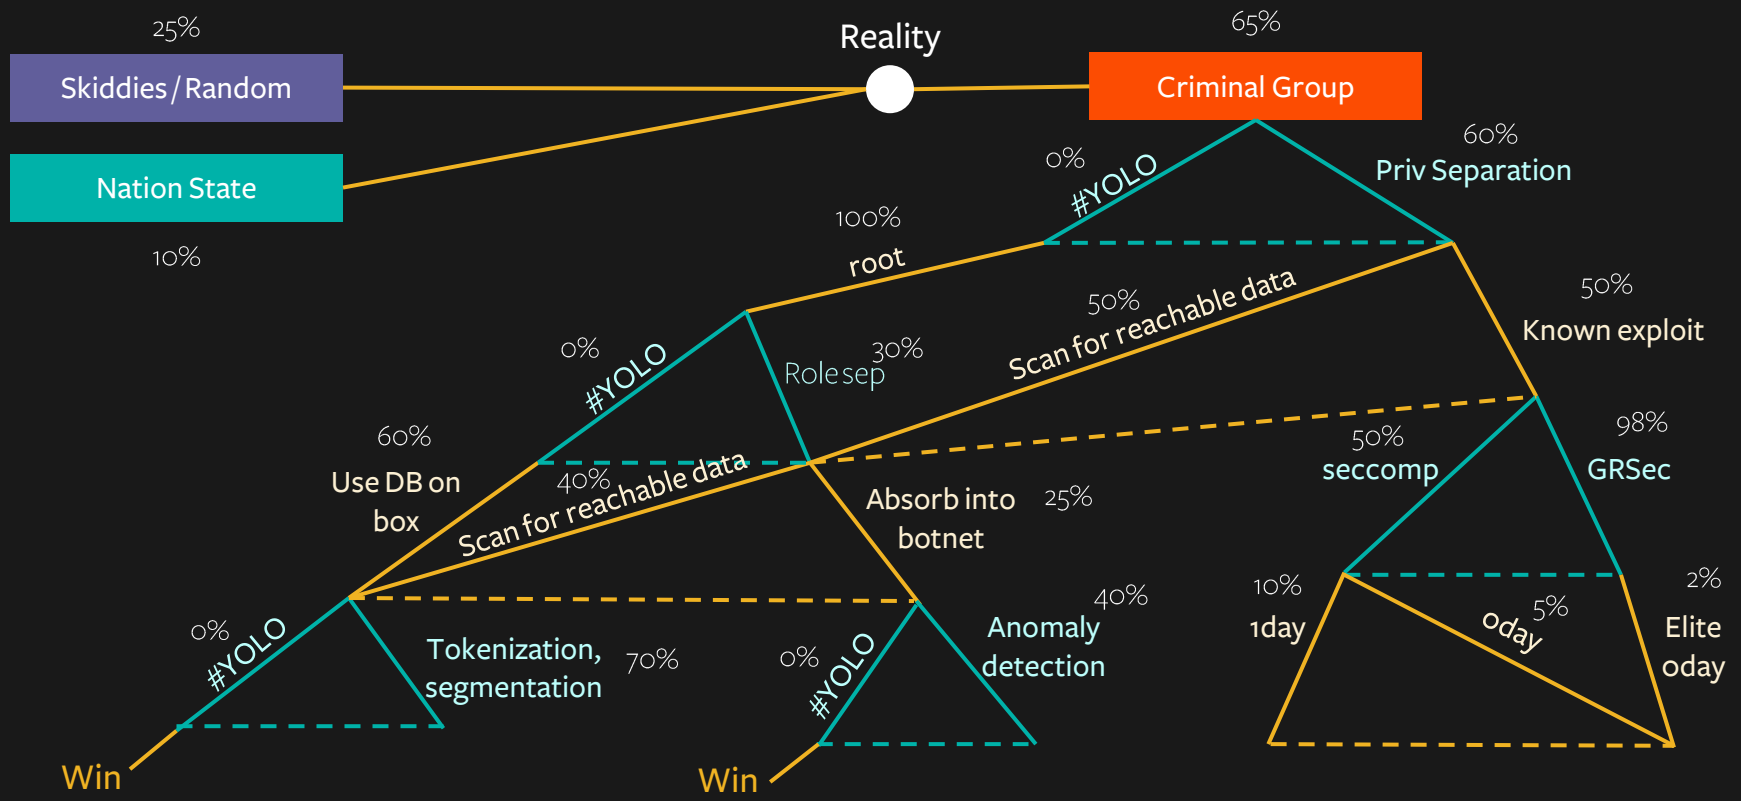

Feedback loop

- Decision trees help for auditing after an incident & easy updating
  - Also helps with general auditing to ensure decisions are revisited and there's not an "additive-only" approach
  - e.g. when info on an attacker group comes out, update the model
- Historical record to refine decision-making process
- Mitigates "doubling down" effect by showing where strategy failed

Decision prioritization

- Defender's advantage = they know the home turf
- Visualize the hardest path for attackers – determine your strategy around how to force them to that path
  - Remember attackers are risk averse!
- Commonalities on trees = which products / strategies mitigate the most risk across various attacks

As a leader of a group

- Leaders shouldn't state biases beforehand
- Solicit feedback that doesn't pressure dissenters to fit majority
- Ask for long-term view of probabilistic costs and benefits
  - Allows room for longer-term projects with high objective benefit
- Get group feedback on decision payoff matrices to compare options – product est. to help X% against attack with Y% likelihood of occurring

As a boss

- Framing is critical – need to add context
- Work with team members to map out probabilities of success or failure of different decisions
- Also, clear ideas of what constitutes success or failure for each decision
- Allow team members to refuse projects without penalty
- Discourage risk taking to “show off” skill level


Ideal decision-making process
(hopefully as cool as this pic)

Bias-resilient process

1. State beliefs about adversaries
2. Model decision trees
3. Spectrum of success / failure for each decision
4. Probability / payoff matrix for different decision options
5. Prioritize rationality over risk-taking
6. Revisit decision trees after each incident


Enlightened Conclusion

Final thoughts

- Make Defense Sexy Again
- Understanding your weaknesses is empowering
- Auditable record of decision process is your best hope
- tl;dr – state assumptions, estimate outcomes (probability & objective benefit), compare with actual results

Further reading

- My upcoming talk at Troopers “Volatile Memory” in March
- My blog post, Behavioral Models of InfoSec https://medium.com/@kshortridge/behavioral-models-of-infosec-prospect-theory-c6bb49902768#.8us8nvycq
- “Two paradigms for depth of strategic reasoning in games” by Zhang & Hedden
- “Skill reputation, prospect theory and regret theory” by Harbaugh

Questions?

- Email: kelly@greywire.net
- Twitter: @swagitda_
- LinkedIn: /kellyshortridge