CIS 350: AI for Cybersecurity

Week 4 Activity Worksheet: Cost-Benefit Trade-Off Under Pressure

CISA-Style Incident Tabletop Exercise (Teams of 3-5)

Team name: _____ **Date:** _____

Team members: _____

_____

Activity Overview

Your team is the security operations center (SOC) for NorthPeak Retail during a Black Friday crisis. An AI model flags suspicious logins, and a flagged account is blocked until verified. You will run THREE timed rounds (about 5 minutes each). In each round you read the situation, weigh the options, and make a team decision that you justify in dollar terms. A complication is injected between rounds. This structure mirrors the CISA Tabletop Exercise Packages (cisa.gov/resources-tools/services/cisa-tabletop-exercise-packages), which real organizations use to practice incident response before a crisis hits.

The Cost Model (use this all activity)

False Positive (block a real customer) = about **$100** (lost sale, support call).
False Negative (miss a real attack) = about **$50,000** (breach cleanup, fraud, reputation).
One missed attack costs as much as 500 wrongly blocked customers. This asymmetry should drive every decision.

Framework: Expected cost = (chance of the error) x (dollar cost of that error). Lower expected cost wins. Accuracy alone is dangerous because it treats a $100 mistake and a $50,000 mistake as if they were the same.

Warm-Up (2 min): Predict, Then Compare

Before the clock starts, each member privately writes a one-word gut answer: which single mistake would hurt NorthPeak most on Black Friday - blocking ONE real customer, or missing ONE real attack? Then compare answers as a team in 30 seconds. Do NOT do the math yet. You will revisit this prediction in the debrief to see whether the cost model changed your mind.

My one-word prediction: _____ **Team split (blocked / missed):** _____ / _____

Task 1: Assign Roles

Assign each role and write the person's name. In a team of 3, one person doubles up on Comms + Legal.

Role	Responsibility	Name
Incident Commander	Runs the clock, forces a decision, owns the call	

Tech Lead	Interprets the model, precision/recall, threshold effects	
Comms Lead	Drafts messages to customers and leadership	
Legal / Compliance	Watches risk acceptance, keeps the audit trail	

Round 1: Set the Alarm (about 5 minutes)

Situation: Alerts are climbing on Black Friday morning. The current threshold catches most attacks but is also blocking a noticeable number of real shoppers.

Decision to make: Keep, tighten, or loosen the detection threshold?

Options to weigh:
- Tighten (block more): fewer missed attacks, more angry customers
- Loosen (block less): happier customers, more risk of a missed attack
- Hold: is 'do nothing' defensible with the numbers you have?

Team decision:

Justification (use dollars): Which error is more expensive here, and why does that make your choice the lower expected-cost option?

Round 2: The Complication (about 5 minutes)

Injected information: Threat intel confirms a real account-takeover campaign is actively targeting retailers like you RIGHT NOW. At the same time, the CMO calls: cart abandonment is spiking because blocked customers are giving up and leaving.

Decision to make: Does this new evidence change your threshold? Re-run the framework. Did the CHANCE of an attack just go up? Did the COST of blocking customers just go up?

Revised team decision:

What changed and why: Did the evidence shift the chance of an attack, the dollar cost, or both? If your decision changed from Round 1, explain exactly what tipped it.

Round 3: Final Call and Documentation (about 5 minutes)

Situation: The morning is ending. Leadership needs a final decision and a written record before the post-incident review. Lock in your final threshold posture, then write the two artifacts below.

Final threshold decision:

Artifact 1 - Message to non-technical leadership (3-4 sentences, plain words, uses dollars, states the risk you are accepting):

Artifact 2 - Risk-acceptance note for the audit file. Include all four: (a) the decision and when it was made, (b) the reason in dollar terms, (c) the specific risk you knowingly accept, (d) who approved it:

Task 5: Debrief (answer together)

First, look back at your Warm-Up prediction on page 1. Did the cost model change anyone's one-word answer? Note who flipped and why.

Q1: What changed your decision between Round 1 and Round 2? Was it the chance of an attack, the dollar cost, or both?

Q2: If leadership asked you to lower the false-positive cost by improving verification, how might that change your ideal threshold?

Q3: Lab 4 connection. In Lab 4 (Phishing Judge) a false positive costs $5 and a false negative costs $10,000, and the optimal threshold turns out to be very permissive (about 0.1 to 0.3). Explain in your own words WHY the huge false-negative cost pushes the optimal threshold so low, and connect that to today's $50,000-vs-$100 model.

Q4: What governance gap did this reveal? Who in a real company should own the authority to change a detection threshold during a live incident, and how should it be recorded?

Grading Rubric (20 points)

Criterion	Description	Points
Completion	Participates in all 3 rounds; every decision is recorded	8
Reasoning	Decisions justified with risk, cost, and evidence (dollars)	7
Communication	Clear stakeholder message and audit-ready documentation	5

Total: 20 points

Scoring: Completion (8) + Reasoning (7) + Communication (5) = 20 points. You are graded on the QUALITY of your reasoning, not on being 'right' the first time.

How to Submit

1. Complete all three rounds and the debrief on this worksheet
2. Record each decision and its dollar-based justification
3. Include the leadership message (Artifact 1) and the risk-acceptance note (Artifact 2)
4. Submit ONE PDF per team with all member names listed on page 1
5. Submit on Canvas by [DUE DATE]

Questions? Post in the course forum or attend office hours!