

Embracing the Kobayashi Maru

Why You Should Teach Your Students to Cheat

Adversaries cheat. The good guys don't. In academic institutions around the world, students understand that they'll be expelled if they violate their college's honor code or otherwise fail to play by institutional rules. But the dissonance between

little advance warning about the exam, and insurrection immediately followed. Why were we giving them such an unfair test? What conceivable purpose would it serve? Now that we had their attention, we informed the class that we didn't expect students to actually memorize the digits of pi, but that we did expect them to cheat.

How they chose to cheat was entirely up to them. Collaborative cheating was also encouraged, but everyone involved would fail the exam if caught. To provide an additional incentive, we offered a prize to the student who exhibited the most creative and effective cheating technique.

The Techniques

Of the 20 students in the course, none were caught, and all of them used diverse approaches. One student used his Mandarin Chinese skills to hide the answers. Another built a small PowerPoint presentation consisting of three slides (an all-black slide, a digits of pi slide, and another all-black slide), the idea being that he could flip to the answer slide when the proctor wasn't looking and easily flip forward or back to a black slide after peeking. Several students hid answers on slips of paper under the keyboards on their desks. One hand-wrote the answer on a blank sheet of paper (in advance) and simply turned it in, exploiting the fact that we didn't pass out a formal exam sheet. Another just

GREGORY
CONTI
*US Military
Academy*

JAMES
CAROLAND
*US Cyber
Command*

how our adversaries operate and how we teach our students puts our students at a distinct disadvantage when faced with real-world opponents who inevitably don't play by the rules. Breaking through the paradigm where students self-censor their ways of thinking to a new archetype that cultivates an effective adversarial mindset is both necessary and possible.

An adversary examines systems and finds weaknesses in trust relationships, human behavior, communications protocols, physical security, and system logic to find exploitable vulnerabilities. By anticipating such actions and reactions, ethical actors are far better prepared to build secure systems and perform both defensive and offensive activities successfully. For both the attacker and the defender, a devious mind is equally as important as a beautiful one.

This article describes our experiences in helping students develop an adversary mindset by adopting the Kobayashi Maru training exercise employed in the fictional *Star Trek* universe. In this exercise, Starfleet cadets face a no-win sce-

nario—attempt to rescue the crew of a disabled civilian vessel and be destroyed in the process, or avoid confrontation and leave the disabled ship and its crew to be captured or killed. Famously, Captain Kirk beat the scenario, and this is important, by stepping outside the game and altering its rules to his benefit. By deciding to cheat and alter the program driving the exercise, he won the contest.

Lest there be any misunderstanding, our purpose here is not to encourage or teach students to cheat in general, but to have them learn to think creatively when considering adversary behavior.

The Challenge

Our variation of the Kobayashi Maru utilized a deliberately unfair exam—write the first 100 digits of pi (3.14159...) from memory—and was part of the pilot offering of a governmental cyber warfare course. The topic of the test itself was somewhat arbitrary; we just needed a scenario that would be too challenging to meet through traditional studying.

By design, we gave the students

memorized the first 10 digits of pi and randomly filled in the rest, assuming the instructors would be too lazy to check every digit. His assumption was correct.

The finalists were particularly innovative. The runner-up used two different techniques, a primary and a backup. In his first approach, he remade his desktop nameplate to look legitimate, but the side facing him included the answer in fine print. For his backup plan, he put the answer on a soda can that he concealed with his hand when the proctor walked by (see Figure 1b). The winner created a false book cover for a course text and replaced portions of the text with the answer, matching text color, font, and size (see Figure 1a). He then used hair spray to lightly tack the false page into place. The result was all but indistinguishable from the original book.

Learning Security Principles from the Cheaters

We learned much from the students during the course of this exercise. They embraced the test, proved far more devious than their day-to-day personas let on, and impressed us with their ability to analyze and defeat the inherently flawed classroom system. We drew the following conclusions from observing the techniques the students used and through an interactive group discussion in which they described their cheating, what they learned, and other techniques they might employ in the future.

The Environment

Students instinctively analyzed their environment and found weaknesses they could use to their advantage. For example, both the presence of computers and the fact that they didn't have to clear their desktops during the exam provided ample opportunity to use their environment to help them cheat.



Figure 1. Examples of student cheats. (a) False book cover containing the answer, and (b) a soda can with an answer sheet that could be concealed when the test proctor was nearby.

Because students were seated side by side and partially hidden behind monitors, some used these characteristics to further facilitate their cheating activities.

Trust

Explicit or implicit trust models are exploitable opportunities. Despite our awareness that the students were cheating, we still inadvertently let our guard down. For example, we wouldn't have stopped a student from using the restroom during the exam. During our group discussion, students suggested that going to the bathroom to cheat would have been an easy-to-implement approach. It's because of our inherent and unconscious trust that we leave ourselves open to exploitation in the physical world and online. As security professionals, we must learn to think like the jaded police officer or prison guard who never takes statements and actions at face value.

Personal Skillsets

Each student possessed diverse skills that he or she could apply to the challenge of cheating. Adversaries do the same. For example, the student who used Mandarin Chinese to write the answer and place it in plain sight used his un-

common skill to become a formidable adversary.

Being Human

Because we're lazy, trusting, and predictable, humans are often the weakest link in any security system, and the students intuitively exploited this fact. One student observed that we rarely handed out worksheets and frequently asked students to provide their own paper. This provided a security gap where they could sneak in an already completed exam and turn it in. Another student suggested instructor predictability and misplaced trust as a potential attack vector. Because we frequently took extra paper from the printer tray to provide to the class, the student said he would pre-position answer keys in the printer and then ask us for a sheet of paper. We would then hand them the answers without knowing it, despite coming from a "trusted" source.

Backup Plans

Adversaries rarely seek to accomplish their objectives through a single, all-or-nothing plan. Several students demonstrated this principle by developing backup plans in case their primary cheating tactic was compromised.

Tips for Teaching Your Students to Cheat

The key to teaching students to cheat is to provide context. Explain to them the objectives of the exercise, which are to learn how an adversary thinks and operates by deliberately loosening traditional rules and tapping personal creativity. While we advocate teaching students to cheat, instructors must still provide clear boundaries, lest there be misunderstandings. In our case, we made it clear that we expected students to cheat and that getting caught would result in a failing grade, but that this exception to traditional rules of behavior only extended to this exercise and not to other graded events in the course.

We deliberately provided minimal warning for the exercise to increase stress levels and material that couldn't be readily learned through traditional studying. During the exam, we sought to further increase the stress and realism by walking occasionally among the student desks. We didn't try all that hard to catch students, but that wasn't the point. We sought merely to increase pressure by acting as realistic exam proctors. We considered but chose not to go as far as forcing students into a position where they must cheat on their own initiative without being told to do so. We believed this would place students into an unfair ethical dilemma, send the wrong message, and most, if not all, students would simply fail the exam rather than cheat illicitly.

Toward a Larger Mindset Curriculum

Our Kobayashi Maru exercise was part of a larger set of lessons designed to cultivate an adversary mindset. There isn't space in this article to describe them in similar depth, but we offer here some highlights to assist educa-

tors in considering more comprehensive approaches.

Early in the course, we included the *Hackers Are People, Too* documentary (Managed Mischief, 2008) to help students understand the hacker mindset, which is sometimes playful and sometimes adversarial. We also included a "divergence" exercise inspired by hacker Dan Kaminsky, who posed the question in the documentary, "What are the alternative uses of a fork?" This seemingly simple question contains significant depth. Students frequently encounter "convergence" questions that seek only a single correct answer. Divergence questions, on the other hand, are open ended and compel students to creatively consider a broad range of answers. We chose this exercise to warm students up to new ways of thinking about problem solving.

Early in the course, we held a lock-picking lab and taught students how to pick small padlocks. The point here, in addition to being a fun, hands-on exercise, was to challenge students' assumptions about physical security and derive commonalities between system security and approaches to understanding and defeating locks.

Our course also included Joe Grand, Jake Appelbaum, and Chris Tarnovsky's case study of insecurities in the San Francisco parking meter system, which taught students how an adversary might attack critical infrastructure (<https://www.defcon.org/html/links/dc-archives/dc-17-archive.html>). For future work, we're considering including a hands-on hardware hacking exercise to teach students how an adversary might develop or modify hardware by building a TV-B-Gone (www.ladyada.net/make/tvb-gone/) universal remote control.

In addition, we included a video by Johnny Long on no-tech hacking to illustrate how an ad-

versary might use social engineering attacks to compromise humans and human-centric security systems (www.defcon.org/html/links/dc-archives/dc-15-archive.html). In the future, we plan to add a phishing email writing contest to give students a hands-on exploration of social engineering.

Weekly throughout the course, students read books to explore various aspects of the adversary mindset, including *Ender's Game* by Orson Scott Card, which illustrated the need to adapt to intelligent adversaries; *Little Brother X* by Cory Doctorow, to teach students the importance of electronic civil liberties and the potential for an adversarial relationship between a government and its citizens; and *Fatal System Error* by Joseph Menn, to examine the real-world actions and reactions between network defenders and online criminals.

Teach yourself and your students to cheat. We've always been taught to color inside the lines, stick to the rules, and never ever cheat, but in seeking cybersecurity, we must drop that mindset. It's difficult to defeat a creative and determined adversary who must find only a single flaw among myriad defensive measures to be successful. We must not tie our hands—and our intellects—at the same time. If we truly wish to create the best possible information security professionals, being able to think like an adversary is an essential skill. Cheating exercises provide long-term remembrance, teach students how to effectively evaluate a system, and motivate them to think imaginatively. Cheating will challenge students' assumptions about security and the trust models they envision. Some will find the process uncomfortable. That's okay, and by design, for it's only by learning the thought processes of our adversar-


ies that we can hope to unleash the creative thinking needed to build the best secure systems, become effective at red teaming and penetration testing, defend against attacks, and conduct ethical hacking activities. □

Acknowledgments

We thank T.J. White and Peiter “Mudge” Zlatko for their feedback and ideas in support of this work. The views in this article are the authors’ and don’t reflect the official policy or position of the United States Military Academy, the Department of the Army, the Department of the Navy, United States Cyber Command, the Department of Defense, or the United States Government.

Gregory Conti is an associate professor in the US Military Academy’s Department of Electrical Engineering and Computer Science and is responsible for the academy’s information security education program. Contact him at gregory.conti@usma.edu.

James Caroland is a member of the US Cyber Command Commander’s Action Group and an adjunct associate professor in University of Maryland University College’s Cybersecurity Program. Contact him at jlcarol@cybercom.mil.

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.