Think Evil®
The Security Mindset

Nicholas Weaver
International Computer Science Institute

Introduction

- This tutorial is largely about how ***I*** think about security problems
  - ➡ And is an attempt for me to understand ***why*** I think the way I do
  - ➡ Your mileage will vary
- This tutorial is very anecdote-centric
  - ➡ They make excellent examples
  - ➡ Exposure to ***many*** different stories is how I learned to think this way

Disclaimers

- Although my research is sponsored by the National Science Foundation, all opinions are my own and not of any funding institution
- Thinking in this way can be, well, bad for your long term mental health:

➡ “The problem is that there is no one arranging meetings where you can stand up and say 'My name is Sam and I'm a really suspicious bastard'”
-Terry Pratchett

The First Story: Casino Cheating

- Most casino games are ***random*** and ***independent***
  - ➡ Winning is based on a ***true random*** process
  - ➡ Both odds of winning and payout are ***independent*** of history
    - ➡ The odds of winning at Roulette: 1 in 38
    - ➡ The payout on winning a $1 bet: $36 (including your $1 back)
      - ➡ Expectation value: 36 * (1/38) = .9474
      - ➡ House advantage: 5.26%
    - ➡ "No one can possibly win at roulette unless he steals money from the table while the croupier isn't looking." — Albert Einstein
- A casino can ***only*** work when the house advantage is positive
  - ➡ Otherwise, it will lose money over time
- Cheaters can ***only*** prosper when they can destroy the house advantage
  - ➡ But cheaters don't have to abide by the rules...
  - ➡ And money is on the line…
- Thus there can be ***no*** agreement between cheaters and casinos: The two groups have goals which are completely opposed

But not all casino games are ***random*** and ***independent***

- Some games have ***history*** on the jackpot size:
  - ➡ Progressive slots and video poker
    - ➡ Jackpot increases with each time any player in the pool does not win
  - ➡ Such games may have positive expectation value for ***individual*** bets
  - ➡ EG, a 3 wheel random slot machine with 30 positions per wheel, and just a single progressive jackpot
    - ➡ Odds of winning the jackpot: 1 in 30^3
    - ➡ Thus if the jackpot is greater than 27,000x the amount bet, the better's expectation value is >1
- Such games ***do not*** affect the house's profit
  - ➡ The growing jackpot is funded by a ***fraction*** of the house's winnings
  - ➡ For the expectation to become positive, a lot of people bet when the expectation was negative

But some games have ***history*** on the odds of winning

- If a deck of cards is ***not*** reshuffled after each hand, subsequent hands are affected by history
  - ➡ Baccarat and Blackjack
  - ➡ E.G: the first hand has 2 aces played:
    Now all subsequent hands will not have these cards
- Players can change their bets based on history
  - ➡ Thus if the odds are favorable, the player can bet more, and can even walk away if the odds become too unfavorable
  - ➡ Which potentially allows a player to gain an advantage over the casino

Edward Thorp Beats The Dealer

- Edward Thorp was a MIT mathematician in the 1960s
  - ➡ Realized both of the previous observations
  - ➡ Realized that with access to a computer (IBM 704) he could develop strategies and run trials
- Developed the basic systems behind ***card counting*** for both Blackjack and Baccarat
  - ➡ Easy ways to track what the deck odds are
  - ➡ Betting strategies to take advantage of shifting odds
  - ➡ Converts a 5% house advantage into a 1% player advantage
    - ➡ This is ***completely intolerable*** for a casino: a casino which allows a player advantage will become bankrupt
- Eventually wrote the book on the subject: ***Beat the Dealer***

Defending against Card Counters: ***Recognition*** and ***Response***

- A card counter must have a predictable pattern of play
  - ➡ Otherwise, the player will not be able to take advantage of when the odds shift in the player's favor
- Thus the Casino can count cards too...
  - ➡ Can distinguish the ***Lucky***
    - ➡ A lucky player is good for the casino: luck doesn't last forever
    - ➡ A lucky player's behavior is not correlated with the state of the deck
  - ➡ From the ***Card Counter***
    - ➡ A solo card counter ***must*** change his bets in response to the state of the deck
- Now simply kick the card counter out of the casino...
  - ➡ Card counting may be ***legal***, but they do not have to let everyone ***play***: A casino can simply kick out a successful card counter

A Defensive Theme: Pattern Recognition

- Many defenses rely on recognizing something as ***good*** or ***bad***
- Anti-virus systems:
  - ➡ Recognize the patterns of ***known*** viruses
    - ➡ Now we can block the bad
- Host-based IDS:
  - ➡ Recognize the behavioral pattern of ***known*** programs
    - ➡ Now we can ***only*** allow the good
- We need to be able to both ***define*** good or bad and ***recognize*** future instances

Defending Against Card Counting: Changing the Rules

- Using more decks makes it harder to card count
  - ➡ More information to keep track of
  - ➡ Odds shifts are considerably smaller
    - ➡ And the odds have to shift greatly for card counting to pay off
- Reshuffle the deck more often
  - ➡ Destroy all history and resets the card counter's job
- Defenses interact synergistically
  - ➡ Reshuffling plus more decks combines to make the problem of counting significantly harder

A Defensive Theme: Change The System/Add Constraints

- In many areas of computer security, the defender controls the rules of the game
  - ➡ The network operator can say what does and does not run on his system
- Changing the basic system can change the entire threat model
  - ➡ Windows XP before Service Pack 2:
    Many network services are ***on by default*** and ***accessable from any remote system***
  - ➡ Windows XP SP2:
    All network services are ***off by default*** and, even when enabled, often ***only accessible from the local network***
- Resulting change has a huge impact on the attack surface: the ways an attacker can compromise a Windows Desktop.

Defense against card counting: Tolerate it (within limits)

- A ***bad*** card counter plays worse than a generic player
  - ➡ A generic player has a -5% advantage,
    if bad play swings this to -6-7%, the house is very happy
    - ➡ And it is easy to mess up
    - ➡ And even a "not quite perfect" card counter might not beat the house advantage:
      As long as their expectation value is still negative, such card counters are ***good*** for the casino
- So unless the card counter is ***winning***, let him continue to count cards!
  - ➡ And if your ratio of bad card counters to good card counters is high enough, just don't bother at all!
  - ➡ The casinos thrive when people ***think*** they can beat the house

Defensive Theme: Tolerance

- Sometimes its not worth putting up defenses:
  - ➡ "You don't put a $10 lock on a $1 rock"
    - ➡ Sometimes the most ***cost effective*** defense is to simply not bother
- Its actually quite common in everyday life
  - ➡ I will personally happily leave a $7 paperback sitting on the table at my local Peet's coffee…
    - ➡ But I will not leave my laptop!

The MIT Card-Counting ring

- A key insight:
  - ➡ Casinos are looking for ***individual*** card counters, but players can join and leave tables at will
- Thus the MIT ring developed ***collaborative*** card counting:
  - ➡ One player at each table plays “basic strategy” (-5% a bet expected return) for low stakes, but keeps track of the count
  - ➡ When the count becomes positives, a “whale” joins the table and bets heavily
- Became the subject of the book ***Bringing Down the House***
  - ➡ This technique is still reportedly in use by other card counting rings: and it is ***legal***

Attacker Theme: Attacking Pattern Recognition

- If you know the defender is looking for particular behavior, adapt your attack accordingly
  - ➡ ***Mimicry***: If the defender is looking for ***known good*** behavior, make your bad behavior look like the good behavior
    - ➡ Mimicry attacks against host-based intrusion detection systems
  - ➡ ***Evasion***: If the defender is looking for ***known bad*** behavior, make your bad behavior look different
    - ➡ Polymorphic/encoded viruses are an evasion attack on antivirus systems
- Mimicry and evasion are common problems with many (but ***not all***) pattern based defenses
  - ➡ The goal is defenses with ***complete*** coverage:
    E.G. If all paths of an attack are covered, ***evasion*** becomes impossible

General Theme: Reaction time

- Part of the reason the MIT ring was so successful was its ***novelty***
  - ➡ The casinos had not expected distributed card counting
- It always takes time to ***react to changes***:
  - ➡ Until the casinos change how they react to the threat, the problem remains
  - ➡ If your opponent has limits on ***adaptability***, exploit them...

Defensive Theme: Cooperation and Communication

- Casinos are not independent, they actually ***cooperate*** on many security problems
  - ➡ Communicate list of people to watch out for
  - ➡ Bulletins about new strategies and tactics
- Once ***one*** casino learns about a new problem, ***all*** casinos may know about the problem…
  - ➡ Even if someone is your competition in ***other*** areas, it often pays to cooperate for security
- Thus the defender is not just an individual entity, but may be a collaboration of multiple entities
  - ➡ Information sharing can be a very powerful defense

But what about Roulette?

- How to steal from the Roulette table:
  After the ball has landed in the slot,
  just change your bet!
  - ➡ Known as "***pastposting***", and represents a major threat
  - ➡ After all, Einstein says you can win at Roulette this way...
- In the early 90s, a company introduced a "no pastposting" roulette table:
  - ➡ An alarm would sound if a player encroached on the play area
- So what is a roulette pastposting gang to do?
  - ➡ Richard Marcus's (a self proclaimed casino cheat's) Solution: Trigger the alarm, repeatedly!
    - ➡ http://www.richardmarcusbooks.com/downloads/19-20%20tech.pdf
    - ➡ The gang members act like drunken idiots, passing items over the table during play
  - ➡ Repeatedly sets off the alarm until the casino pit boss just turns it off
  - ➡ Once the alarm is turned off, ***then*** steal the table blind...

Attacker Theme: Malicious False Positives

- If a defense is triggered when there is no attack, this is a ***false positive***
  - ➡ If an attacker can trigger this, you have ***malicious false positives***
- Many uses for malicious false positives
  - ➡ Get a system deactivated due to frustraction
  - ➡ Distract attention from the real target
  - ➡ Cause damage due to the defenses themselves
    - ➡ Reactions have a cost: the attacker may simply wish to cause the defender to face these costs

But why not just ***corrupt*** the dealer?

- If you're "friends" with the dealer, who says things have to be random?
- Many ways for a corrupt dealer to cooperate with an accomplice:
  - ➡ On Blackjack: The dealer needs to check if he got a blackjack when an ace is showing
    Dealer behavior can signal this to an accomplice at the table
    - ➡ Accomplice can then do an "insurance" bet
  - ➡ On Blackjack or Baccarat: The dealer switches the deck with a prepared deck
    - ➡ This attack can be ***deadly***, as a whole table of accomplices gets incredibly "lucky": Very high risk but very high reward
  - ➡ On just about any game: Just be "stupid": miss cheating attempts such as switching chips, late bets, or other behavior

Attacker Theme: Insider Attacks

- The ***insider attack*** is often the most insidious:
  - ➡ Insiders ***must*** be trusted, the attack is a betrayal of trust
  - ➡ Insiders ***must*** have detailed knowledge of the system
  - ➡ Insiders are ***people***, with all the human weaknesses
- Casino cameras have to watch the dealers as much as the customers
- Why do you think Costco, Fry's, etc check receipts at the door?
  - ➡ Its to prevent a cashier from colluding with a customer to sell a big-screen TV as a can of Coke...

But Roulette Tables are Getting Smarter

- Some casinos are experimenting with RFID (Radio Frequency ID) casino chips
  - ➡ Each chip has a unique serial number and RFID chip
  - ➡ The table can use this to monitor where ***every*** chip is on the board
    - ➡ Can also monitor, in real time, who has what
    - ➡ Also makes forging chips considerably more difficult
- Now the roulette table can directly detect pastposting
  - ➡ By keeping track of when each chip is added or removed
  - ➡ Can detect otherwise very hard to detect moves
    - ➡ Such as placing a stack of chips with an almost hidden high-value chip, which is swapped out with a only low-value stack on failure
      - ➡ Because most croupiers and cameras are looking for people adding bets to winners, not switching bets on losers

But Einstein Was Wrong, You CAN Win At Roulette...

- Thorp also observed that Roulette is ***not*** a random process...
  - ➡ IF AND ONLY IF bets are allowed after the ball is spun on the wheel
- Collaborated with Claude Shannon to develop a Roulette-tracking wearable analog computer in 1961
  - ➡ Toe switch to input data, earphone for output
- Idea:
  - ➡ Track the velocity and phase of the rotor and ball
    - ➡ Measure by clicking the switch when the rotor and ball pass certain points
  - ➡ Tone indicates what octant the ball was most likely to hit
- Amazingly effective: >40% player advantage in both the lab and the casino!
- "The Invention of the First Wearable Computer", E. O. Thorp

The Casino Responses...

- "Place your bets" and ***then*** spin the ball
  - ➡ Restores randomness to the game, if the casino does this
    - ➡ Not all casinos do: there was a case in 2004 where this technique was employed using a laser-scanner "cellphone"
      http://www.guardian.co.uk/science/2004/mar/23/sciencenews.crime
- Change the law:
  - ➡ In 1985 (when such devices were becoming far more common) an emergency measure was passed in Nevada:
    Using a technological device to aid in gambling is a felony
    - ➡ Target was not just roulette computers but easy-to-use Blackjack counting computers
- This grossly changes the stakes for cheaters who get caught
  - ➡ Many more potential "cheaters" still want to operate within the law:
    It changes the ***costs*** involved in cheating

Defensive Theme: Change the attacker's costs

- Attackers have many costs in their attack...
  - ➡ Not just the cost of the attack, but the cost of being caught factored into the probability of being caught
- Anything which changes the attacker's cost model may dissuade attackers
- Also there is a "Bear Race" factor
  - ➡ "I don't need to outrun a bear, I just need to outrun the guy I'm standing next to."
    Make attacking ***you*** more difficult than attacking your neighbor

Outline

- People, Ideas, and Technology...
  - ➡ Strategy and Tactics
  - ➡ Adversarial modeling
    - ➡ Informal
    - ➡ Formal (OODA loops)
    - ➡ Attacking decision cycles
- Constraints & End States
- Applications:
  - ➡ Internet Service Providers vs Peer to Peer systems
  - ➡ Worms, Viruses, and Things that go Bump on the Net
  - ➡ Personal protocols to protect my finances
  - ➡ Why High Finance ***must*** fail

“People, Ideas, and Technology... In That Order”

- Col. John Boyd, US Air Force
  - ➡ Developed the Energy/Maneuverability theory
    - ➡ The mathematics behind fighter-aircraft operation: Provides a single-graphic view of airplane performance based on how quickly it can add and dump energy in maneuvers: based largely on thrust/weight and drag
  - ➡ The leading force behind the F16 and F18
  - ➡ Developed much of the modern military theory of conflicts: the ***OODA*** loop process
  - ➡ This is one of his more famous quotations
- In all the casino examples, technology was an ***enabler***, but it was human behavior that is key

So some thoughts about people

- People are self-interested
  - ➡ They usually act in what ***they perceive*** as their self interest
- People are motivated
  - ➡ They know what they want and they will ***try*** to get it
- People are adversarial
  - ➡ When self interests collide, you get a conflict

People are ***self interested***

- People ***usually*** act in their self interest
  - ➡ ***If*** they understand their self interest
- "The universe runs on a mix of energy, matter, and enlightened self interest"
  -G'kar
  - ➡ Of course, every individual's definition of ***self interest*** may be different:
    EG, my primary self interest is to ***enjoy what I do***
    Somebody's interest on wall street may be to ***make lots of money***
- Your opponent's self interest often dictate their strategy and tactics
  - ➡ EG, the authors of malcode for profit (interested in money) behave very differently from the authors of malcode for espionage (interested in information belonging to specific parties)

People are ***motivated***

- Self interest leads to ***motivation***:
  Once people know what they want, they will try to get it
  - ➡ Within the constraints which they will operate
  - ➡ Level of motivation varies:
    I'm rather lazy: do the ***minimum*** needed to accomplish my objectives
    Others may be ambitious
- Once you understand the participants self interest and level of motivation, their objectives should be (reasonably) clear

People are ***adversarial***

- If ***your*** self interest opposes someone else's self interest, competition may be inevitable
  - ➡ It can be subtle and normal:
    Economics is all about competetion and adversarial behavior
  - ➡ It can be overt and illegal:
    Criminal adversaries
- This ***creates*** conflicts:
  - ➡ Do the different parties have different self interests?
- This can also ***diffuse*** conflicts:
  - ➡ Can you change the system so that different parties' self interest ***aligns***?

Strategy

- The high level techniques for accomplishing a particular goal
  - ➡ The high-level ***Why*** of the conflict
- This is usually centered around the interest of the parties:
  What is their overall objective?
  What is their level of motivation?

Tactics

- The detailed techniques and tools needed to accomplish a local goal

  ➡ The low level ***how*** of the conflict

- This is centered around the motivation of the participants:
  How to achieve the actual objective

Some thoughts….

- We need to work on problems on ***both levels***
  - ➡ Tactics require a strategy to be useful
  - ➡ Strategies require tactics to implement
- Often, strategy is an effective lever
  - ➡ Disrupt the ***why*** of the conflict:
    What is the other guy's interest and objectives?
    Can we change how these operate?

Adversarial Decision Making

- The goal is to not ***beat*** your opponent but to ***drive him insane***:
  - ➡ If you can beat his decision making you should win
  - ➡ For an organization, destroy their decision making process
- So how can we model adversarial decision making?
  - ➡ If we want to realistically attack the opponent's decision making, we need to reason about it clearly
- The informal model:
  Your Evil Twin
- The formal model:
  The OODA loop, developed by Col. John Boyd

The Less Formal Model: Your Evil Twin

- You need to model an adversary who's as resourceful, as creative, and as innovative as possible
  - ➡ You can dumb down your opponent later should conditions warrant, but it is best to assume an opponent who is too ***smart*** rather than too ***stupid***
- If you can model an adversary who is ***more*** resourceful, creative, and innovative then you are…
  - ➡ Simply run that model and ***become*** that person
- Thus the most sophisticated adversary you can actually model ***is you*** (or your evil twin)

Understand Strategic Objectives ***First***

- What is the ***interest*** and ***objectives*** of your evil twin?
  - ➡ Or various evil twins
    - ➡ ***Empathize*** with your evil twins...
  - ➡ Does he want to make money?
  - ➡ Does he want to learn ***your*** secrets?
  - ➡ Does he just want to see the world burn?
- What role does he play?
  - ➡ Is he an Internet Service Provider?
  - ➡ Part of a criminal conspiracy?
  - ➡ Working for the Chinese government?

Then Define ***Resources*** and ***Constraints***

- How much resources does your evil twin have?
  - ➡ Lone wolf: Your evil twin in the basement
  - ➡ Criminal syndicate: A support network of some money and others
  - ➡ Nation-state employee: A full ***clone army*** of evil twins with millions of dollars in backing
- Are there particular constraints?
  - ➡ Does he need to obey the law?
  - ➡ Does he need to worry about public opinion?

Only ***then*** focus on tactics

- The other constraints can define what tactics may or may not be acceptable:
  - ➡ If your evil twin's objectives require legal behavior, only some tactics are in play
  - ➡ Available tactics also depend on the position/abilities
- Its a mistake to rathole on tactics too early
  - ➡ Why try to defend something which the probable attackers wouldn't care to do?

The F-86 vs MiG 15 puzzle and the origins of the OODA Loop

- Boyd was an F-86 Saber pilot during the Korean War
  - ➡ In the Korean war, the F-86 proved superior to the Russian MiG 15
    - ➡ A claimed 10:1 kill ratio!

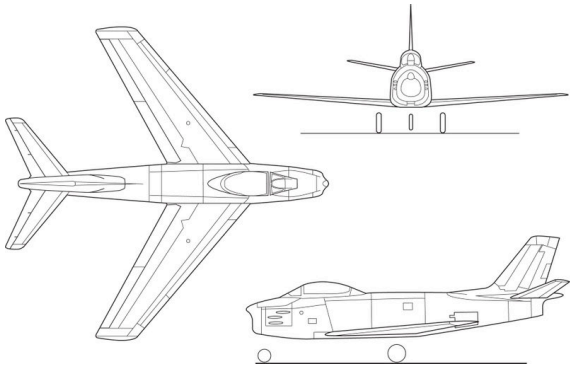

- But based on the physics of the aircraft, the MiG is far superior!
  - ➡ 20% better! thrust/weight ratio
- So why did the F-86 do so well?
  - ➡ Part was doctrine and training:
    US pilots were more experienced and used better air-to-air tactics
  - ➡ But a major factor was “user interface”

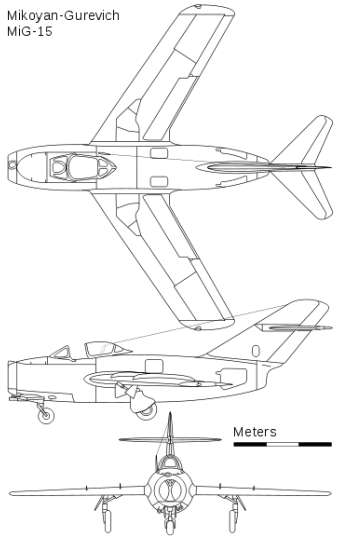

The F-86 and MiG 15 User Interfaces

- Partially, the MiG had an inferior canopy
  - ➡ Less visibility: The pilot sat lower in the plane and there were more obstructions in the view
- But a very big factor: the MiG's controls were inferior
  - ➡ The F-86 used hydraulic controls:
    a light pilot input is sufficient to turn the plane
  - ➡ The MiG 15 used cables:
    The pilot provided all the force needed to move the control surfaces
- Thus in a turning fight, the F-86 had a huge advantage:
  - ➡ The ***pilot*** required less effort to make a ***series*** of maneuvers

Agility

- The F-86 had a huge advantage in pilot ***agility***
  - ➡ The ability to rapidly ***change*** behavior/positions/tactics in response to how a dogfight unfolded
  - ➡ This compensated for worse agility on the plane
    - ➡ Because the MiG had vastly more thrust, it could regain lost energy much faster
- Boyd developed agility-centric air-to-air tactics as a flight instructor
  - ➡ But how to apply this to general decision making?
  - ➡ How do people and organizations come to a decision?

John Boyd's Insights...

- Key insights:
  - ➡ There are multiple processes in creating a decision
  - ➡ Decision making composes:
    You can treat an organization as a single entity composed of individual entities
  - ➡ There are ***fast paths*** and ***slow paths*** in the decision making process
  - ➡ If you can develop a model of how opponents think, you can then ***attack the decision making cycle directly***
- Developed the "OODA Loop" (Observe, Orient, Decide, Act) to describe the decision making process
- Developed a theory of "Moral Conflict" on how to attack the decision making process

The OODA Loop

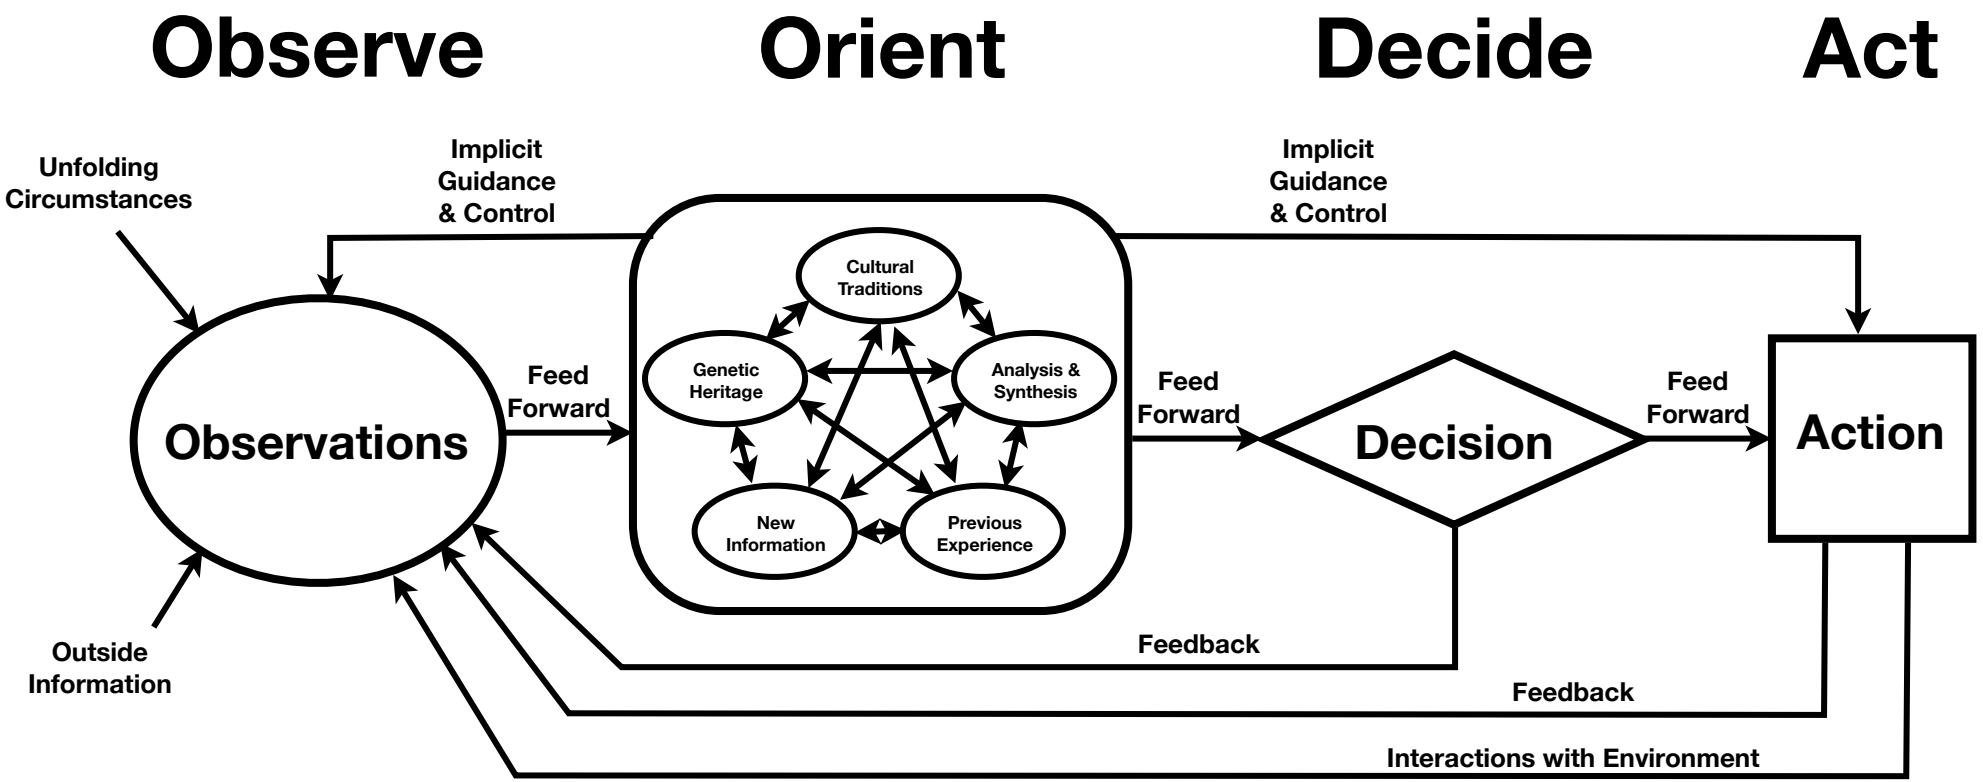

This drawing is based on one from defense and the national interest (www.d-n-i.net)

Observation

- Observations are the data input into the system

➡ How do you collect information on what's going on?

- Eyes, sensors, messages, etc…
- Without ***accurate*** and ***timely*** observation, the decision cycle turns inward

➡ And ends up breaking...

Orientation

- Orientation describes the ***implicit*** decision making processes
- For an individual
  - ➡ Instinct, reaction, training, history, the "snap decisions" which occur all the time
- For a composite entity
  - ➡ The ***decentralized*** decision-making process, where individuals don't need to cooperate or coordinate, but just ***do***
- Orientation is the ***fast*** path

Decision

- The decision process is the ***explicit*** decision making
  - ➡ If you have to go “I need to think about it...”, its an explicit decision
  - ➡ In an organization, this usually involves consensus or discussion
- Explicit decisions are the ***slow*** path
  - ➡ As soon as you have to make an explicit decision, things grind to a halt
    - ➡ Especially true in organizations

Action

- Actually ***do*** something…
  - ➡ In many ways, the ***least*** interesting part of the decision cycle is actually doing anything
- There are plenty of constraints on actions, which we will discuss later...

Communication

- In Boyd's model, communication is an ***implicit*** set of connections
  - ➡ Its an action to send messages
  - ➡ Its an observation which receives messages
- Its often best to think of communication ***explicitly***:
  - ➡ Messages take time, they may be unreliable, etc...
    - ➡ Properties of the communication medium and the organization
  - ➡ Communication media are ***resources*** as well

Automation

- Automated decision cycles occur all the time
  - ➡ Worms and malcode are automated attacks which require automated defenses
- Automation is an even ***faster*** decision cycle
  - ➡ Anything involving human vs automation:
    Automation wins the race.
  - ➡ Distinction between “Orientation” and “Decision” disappears:
    - ➡ There are no clear cut “Fast path”/”Slow path” distinctions for automated decision systems
- Some security problems can ***only*** be addressed with automated decision cycles:
  - ➡ If the attack is automated (eg, a worm),
    reactive defenses (those that detect and respond) must also be automatic

Composition

- An ***organization*** or suborganization has its own meta-OODA-loop
  - ➡ Representing the organization's overall decision making
- Much work in how to construct these organization ***well***:
  - ➡ Delegation of responsibility and "mission based orders": Each individual and small group knows both their individual task and overall objectives
    - ➡ Minimize the amount of ***explicit*** decision making needed
- ***Many*** ways to construct these ***poorly***:
  - ➡ ***Micromanagers*** eliminate delegation of responsibility and add needless explicit decision making
  - ➡ The "Yes Men" phenomenon: ***Incestual Amplification***

Attacking Decision Cycles

- If you can drive your opponent mad, you will win
  - ➡ So the goal is to ***disrupt their decision cycle***: Make it so they can't come to the right decision or, better yet, a decision at all!
- But at the same time, we must strengthen our own process
  - ➡ Otherwise, of course, we might lose

What Boyd wants to create in the opponent...

- The foundation of ***moral*** conflict:
  - ➡ Conflict waged ***directly on the decision making process of the opponent***
- Menace:
  - ➡ The impression of danger to one's well-being and survival
- Mistrust:
  - ➡ An atmosphere of doubt and suspicion
- Uncertanty:
  - ➡ Events that appear ambiguous, erratic, unfamiliar, etc…
- Causes opposing individuals and groups to become ***non cooperative***:
  - ➡ Breaks down their operational structure by increasing ***friction***

But at the same time, strengthen our own institutions...

- Counter menace with ***initiative***
  - ➡ The ability to take action without being urged and under stressful conditions
- Counter mistrust with ***harmony***:
  - ➡ Create friends and influence people
- Counter uncertanty with ***adaptability***:
  - ➡ You can't always be ***certain***, but you can always be ***adaptable***

Speed

- Now with those themes...
- If all else is equal, a faster decision cycle is more adaptable
  - ➡ Can react to unfolding circumstances more quickly
    - ➡ Increases adapability
  - ➡ Can ***create*** unfolding circumstances more quickly
    - ➡ Creates uncertanty in the opposition
- But this is ***only*** if all else is equal
  - ➡ Getting the wrong answer fast still gets you just the wrong answer...

Accuracy

- We still need to get the right answer
- Often the key is observation and data input:
  - ➡ Without accurate data, how can you reach an accurate concusions?
    - ➡ Forces a large emphasis on data collection and data analysis
- Thus a common theme of deception:
  - ➡ Its a direct attack on the opponent's accuracy

Constraints and End States

- Competitions are ***not*** freeform: they occur within constraints
  - ➡ The ability of competitors depends on the constraints in the system
  - ➡ Constraints act to limit the possible actions
  - ➡ Money is a constraint
- Constraints and technologies can drive portions of a conflict to ***end states***:
  - ➡ A tactic or technique will become ineffective to the point of uselessness

Constraints

- Constraints limit the freedom of action
  - ➡ You can't just do ***anything***, it must be actually ***possible***
- Understanding the constraints is necessary to understand the form of the competition
  - ➡ Constraints act as a limiter of possibilities
- But understand that some asserted constraints aren't:
  - ➡ EG, what happens if an attacker can break into your machine room?

Constraints of Physics

- The most basic constraint:
  The speed of light
  - ➡ A related limit: the speed of communication on a network
    - ➡ Acts to limit some defenses
- But systems have “Physics” too:
  - ➡ E.G. an ARM processor runs ARM binaries: you can’t run x86 code on an ARM without an emulator
    - ➡ Therefore, attacks which rely on x86-specific behavior won’t work on an ARM platform

Constraints of Law and Public Opinion

- Some adversaries ***must*** work within the law
  - ➡ Businesses need to be legal to survive
  - ➡ Others strictly do not care, such as criminals
- Public opinion may be as important as legal constraints
  - ➡ Negative public opinion can result in draconian legal restraints being added
  - ➡ Negative PR can be costly on its own merits
- Your adversary's position dictates available actions

Adding Constraints

- Constraints are a great ***preventative*** defense
  - ➡ If you control the environment, constraints can limit an attacker's actions
    - ➡ Firewalls, software restrictions, etc...
- A significant area for usability research: ***Adding*** constraints to a system
  - ➡ Is it possible to make a system such that users ***can not*** perform exploitable actions, yet still be usable?

Cost is a Constraint

- ***Nobody*** has unlimited resources:
  - ➡ Time, money, people, opportunity, all are limited
- All parties must spend their resources ***wisely***
  - ➡ It often includes knowing when ***not*** to secure something:
    "You don't put a $10 lock on a $1 rock"
    After all, who would steal the rock ***10 times***?
- Money is a good metric for the other resources
  - ➡ "Time is money" has a corollary: "Money buys time…"

Consider the money limit...

- When Steve Trimberger at Xilinx designed their FPGA copy protection…
  Targeted adversaries ***only*** willing to spend less than $100,000 to copy a design in an FPGA
  - ➡ Steve assumed it would cost roughly $100K to bribe an engineer to get a copy of the design
  - ➡ So why devote resources to technical defenses which might require $500K to defeat, but could be evaded with a $100K bribe?
- Likewise, a ***captcha*** is ***not*** about determining that someone is human…
  - ➡ Rather it is a way of determining that the user is human ***or*** the user is willing to spend $.0025 or so to ***appear*** human
  - ➡ Thus a ***captcha*** can ***only*** defend a resource which is worth less than that to an attacker
    - ➡ It works to keep blog spam off of most blogs
      - ➡ After all, the Google Adwords are less expensive
    - ➡ It doesn't work to stop scalpers at TicketMaster

End States:

- Not all tactics work forever:
  - ➡ We no longer see mounted knights across the field of battle
- Technologies evolve and can result in the extinction of tactics:
  - ➡ Evolving the system towards an ***end state***
- End states often favor one side or the other

Complete End States

- Some lines of tactics can evolve to a ***complete end state***:
  - ➡ Attacker or defender is effectively stuck unless the basic technology somehow changes
- Happens quite often:
  - ➡ E.G. buffer overflow attacks can't be used against programs written in bounds-checked languages
- If an end state favors ***you***:
  push development towards the end-state
- If an end state favors the adversary:
  Don't fight this battle...

Effective End States

- "Malcode Wars are not won by solving the halting problem.
  Malcode wars are won by making the other poor bastard solve the halting problem."
  - ➡ With apologies to George S Patton
- Some end states are ***not*** guarenteed, but are so much harder for one side or the other:
  - ➡ Unfortunately for Symantec, "Virus Detection" ***is*** the halting problem, thus signature-based detection of malicious code (classic antivirus software) is a losing battle
    - ➡ Thus I'm not interested in writing AV software, or tools to automatically analyze Javascript to detect potential malice

The Rest of the Tutural: Applications and Case Studies

- ISPs, Content Providers, and Peer to Peer Technology
- Worms, Viruses, Bots, and Things that go Bump on the Net
- How I Protect My Wallet
- Why Wall Street Can't Work

A Current Conflict: ISPs vs Content Providers vs Customers

- The ISPs objective:
  - ➡ Provide an acceptable level of service to the customer while maximizing the ISP's profit
    - ➡ Commonly in a ***duopoly***: there exists competition between ISPs but it is somewhat limited
    - ➡ Also constrained by other businesses:
      Almost every major ISP is either a telco that wants to be a cable company, or a cable company that wants to be a telco
- The customer's objective:
  - ➡ To get the ***desired content*** in a way which minimizes the customer's cost
    - ➡ Convenience, hastle, and legal risk are all costs
    - ➡ Different customers weight these costs differently
- The content provider's objective:
  - ➡ To provide ***paying*** customers with their desired content at the minimum cost ***to the content provider***
  - ➡ To limit the impact of piracy

My Perspective In This

- As a researcher developing tools to ***understand*** ISP behavior
  - ➡ Detecting P2P interference using injected RST packets
  - ➡ Developing tools to probe large suites of behavior
    - ➡ Obligatory plug: http://netalyzr.icsi.berkeley.edu
- As a ***rational***, lower-bandwidth customer
  - ➡ I pay for the high bandwidth service, but I'm rather low in GB/month
- As someone who wants to ***diffuse*** potential conflicts
  - ➡ Driven from the ***informal*** model: ***empathize*** with all sides
  - ➡ Develop technologies that can ***unite*** competitor's strategic goals

Round 0:
Old School FTP Warez

- Warez (pirated PC games) were spread on FTP sites and dialup bulletin boards in the 1990s
  - ➡ If you knew the login, you could get the pirated games
    - ➡ EG, a friend's mother's university account supposidly hosted an FTP site with a GB of pirated games...
- Bandwidth issues:
  - ➡ This was not cheap, and sources were identifiable
    - ➡ Criminal charges were filed against participants
- Of consequence, only a limited # of participants
  - ➡ ***Closed world*** piracy: small communities of pirates
    - ➡ Annoying but tolerable for content providers: Limits the number of participants

Round 1:
The Rise of Napster...

- Napster was a peer-to-peer program for sharing music files
  - ➡ Users could make public their folder full of MP3s and copy them between users
  - ➡ Became available in 1999
- Addressed a huge gap in the available software:
  - ➡ Made it much ***easier*** to find and obtain music online
- Enabled ***open world*** piracy:
  - ➡ Rather than having to know ***someone*** to get a piece of content, allows ***arbitrary*** users to find pirated content
  - ➡ But limited to ***small*** files: music rather than video content
- Open-world piracy is much more dangerous:
  - ➡ Lowers the barrier to entry for those wanting pirated content
  - ➡ Lower barrier to entry means that otherwise paying customers may become pirates

The Content Provider's Response: The Court System...

- Content providers sued Napster for contributing/ enabling copyright infringement on the part of the users…
  - ➡ And won, bigtime
- Napster effectively shut down when an injunction was granted which ***required*** Napster to prevent the sharing of copyrighted files
  - ➡ Since Napster controlled the index for the files as well as the software
  - ➡ Napster concluded "We can't do copyright enforcement, so we'll just shut down our index servers…"
- But it was too late…
  - ➡ Customers grew to like open world piracy
  - ➡ Which meant new P2P software could try to profit from it...

The P2P Software Provider's Response: KaZaA

- Several software providers observed there is (potential) profit in building P2P software
  - ➡ Usually by bundling sleazy adware and similar items in the software package
  - ➡ But the Napster lawsuit showed that such software is unlikely to survive a court challenge
- Solution: Incorporation-shopping
  - ➡ Sharman Networks incorporated in Vanuatu! and headquartered in Australia
  - ➡ If you can't be sued if you don't have a business presence (hopefully)
    - ➡ This is an example of ***jurisdictional arbitrage***: taking advantage of differences in law rather than differences in price
- Eventually failed, but the decision cycle is ***slow***
  - ➡ In many ways, ***both*** sides lost:
    Sharman didn't make the pot of gold
    The RIAA didn't stop the piracy

The rise of iTunes Store 2003

- A ***new*** content provider model:
  - ➡ The record labels would not agree amongst themselves to sell music online
  - ➡ But they ***were*** willing to license content to Apple and others
- Created a new set of ***intermediaries***
  - ➡ Unfortunately for the RIAA, it allowed these intermediaries to become more powerful
- Showed that ***paid content*** can be profitable
  - ➡ DRM used to reduce uncertanty
    - ➡ But made the content less desirable to many users
  - ➡ Many users view illegality as a ***cost***:
    Legal content ***can*** compete with illegal content
    If you lower the total ***costs*** of getting legal content
    - ➡ Ease of use is a cost

The best content provider response to date: ***Hulu***, ***Netflix***, etc...

- It is critical to ***not*** just make infringing content more difficult to obtain, but to provide a ***legal alternative***
  - ➡ Users ***will*** pay for content: either directly (iTunes, Netflix Streaming) or through advertisements (Hulu)
- But this is ***not free***: Content delivery costs money
  - ➡ $.10/GB for Amazon
  - ➡ ~$.20/GB? More? for Akamai
    - ➡ Akamai provides a lower-latency service
- And it all adds up:
  - ➡ 1 hour HD video == $.10... Times ten million views…
  - ➡ Credible estimates suggest that YouTube costs $300M a year in bandwidth bills...

The Rise of BitTorrent

- BitTorrent's developers realized the ***real*** weakness of Napster and Kazaa: the ***combination*** of content deliver and seach
  - ➡ Search engines alone are legal
  - ➡ Content delivery protocols are legal
  - ➡ So focus solely on ***content delivery***
- BitTorrent optimized for delivering ***large*** files:
  - ➡ A ***tracker*** keeps an index of who is participating in a ***swarm*** of peers
  - ➡ Individual peers keep track of which content other peers have
- Splits responsibility:
  - ➡ Splits content discovery (the Pirate Bay) from content delivery
- Removes the bandwidth costs:
  - ➡ ***Shifts*** the delivery cost from the (pirated data) provider to the recipients
- Provides a significant non-infringing use for legitimate content providers

Bulk Data P2P
As Cost Shifting

- Bulk Data P2P (BitTorrent etc) offers a way of ***shifting*** the cost of content delivery from the content provider to the content recipients
- ***Necessary*** for piracy of large files:
  Individual users lack the bandwidth or the money
  - ➡ Enables ***open world*** piracy of ***large*** files:
    Without this, anyone who tries to share a large file is going to see crippling bandwidth bills:
    How many pirates will pay $10,000 to share 1GB with 100K other pirates?
- ***Useful*** for legitimate content providers:
  Allows the content provider to shift the costs to the recipients of the data
  - ➡ CNN sees a 30% reduction in bandwidth costs for their P2P streaming browser plug-in vs conventional content delivery
- BitTorrent's basic idea can be adapted to streaming video:
  - ➡ The key observations are that:
    The blocks of transfer should be individual frames and small groups of frames
    The most-desirable blocks for a client depend on where it is in the video stream

But now ISPs Grew Concerned

- Music files are relatively small:
  Even a really REALLY committed pirate will only transfer a few GB
  - ➡ But video files are ***huge***: A single hour may be a GB or more
- And ISP bandwidth cost is considerably more than content provider bandwidth cost
  - ➡ It always costs more to bring 1 Mbps to 100 places than 100 Mbps to 1 place
- And ISPs were seeing congestion effects in their network
  - ➡ Comcast was ***falsely*** accused of disrupting Vonage's VoIP service. The real cause was probably bittorrent-related congestion on the shared cable-modem uplinks

Thus P2P is Inefficient Cost ***Shifting***, not Cost ***Saving***

- Bulk-data P2P doesn't reduce the amount of data transferred
  - ➡ One copy downloaded, one copy ***uploaded*** per customer
- But ISP bandwidth ***must*** cost more than content provider bandwidth:
  - ➡ Content providers can be optimally located
  - ➡ ISPs are ***always*** suboptimally located
  - ➡ And its far less expensive to bring 100 Mbps to one location than 1 Mbps to 100 locations...
- Cable systems in particular are vulnerable:
  - ➡ For them, the last mile is ***very*** costly: Every Mbps to a customer represents bandwidth that could be used for a TV channel
    - ➡ DOCSIS actually encodes the downstream data in MPEG "frames"
  - ➡ The uplink is ***more*** costly because its less efficient
    - ➡ A few P2P users can clog a neighborhood with long lived flows
  - ➡ Thus a perfect P2P system will cost a cable ISP significantly more than a normal download
    - ➡ And because ISP bandwidth is ***more*** expensive, the aggregate costs are substantially magnified

ISP Reaction #1: Managing P2P traffic

- Bulk data P2P is trivial to recognize:
  - ➡ Its just that, ***bulk data*** and ***peer-to-peer***:
    No data cloaking can remove data or remove peers
    - ➡ Traffic analysis: Knowing "Who talks to who", "for how long", "how much data", "what are the patterns in the data" is a very ***very*** powerful tool
- Once you recognize it, now ***do something about it***
  - ➡ ***Block*** some connections:
    Can limit total traffic flow without blocking P2P completely
    - ➡ Blocking can be done using injected TCP Reset packets
    - ➡ Often it is best to block some ***types*** of connections:
      Blocking "seeding" (uploading only) is actually beneficial to most customers, as they usually aren't benefiting
  - ➡ ***Slow*** some connections:
    Put all P2P connections through a synthetic bottleneck and let them fight it out themselves

The Public Reation Was Vicious...

- "Comcast is forging packets!"
  - ➡ Public reaction was more vicious than the ISPs expected
    - ➡ Especially since it was done to prevent other public problems!
  - ➡ Comcast PR didn't help: initial denials and false statements clouded the issue
    - ➡ Their actual implementation was actually pretty good: it ***only*** blocked pure upload flows
- Worse, for the ISP, blocking traffic is detectable:
  - ➡ The Glasnost project had a web site who's applet behaved like P2P traffic to check for blocking
  - ➡ I and colleagues developed detectors for RST injection: able to distinguish injected reset packets (used to block flows) from normal RST packets
- ***Transparency*** limits corporate behavior:
  - ➡ Even if its legal, they don't necessarily want to get ***caught***
  - ➡ Public opinion is now viewed as a huge constraint on ISP actions

So the coming conflict... P2P management

- ISPs are constrained by public opinion:
  - ➡ Can piss off a few activists, but not everybody
- Content providers are constrained by cost:
  - ➡ Your competitors who go with P2P will be spending less money to serve the same number of customers
- Some ISPs are constrained by available bandwidth:
  - ➡ Even with DOCSIS 3, cable ISPs have really limited uplink bandwidth, even with perfect localization
  - ➡ Wireless ISPs are really ***really*** expensive bandwidth

Exacerbating the Conflict: Caps

- Some landline ISPs have proposed low bandwidth caps (~50GB or less a month)
  - ➡ Such caps are significantly anticompetitive:
    Prevents video-on-demand services from being used
- Wireless ISPs almost inevitably have usage limits of ~5 GB with expensive overage charges:
  - ➡ Overage charges add ***uncertanty***:
    One of the little secrets of the iPhone:
    the ***data*** portion is very consistant and predictable
    the ***voice*** portion is reasonable
    - ➡ Data is far less predictable for the average user
- I'm happy with high caps (>250 GB), but lower bandwidth caps (~50 GB or less) I view as a huge danger to Internet innovation:
  - ➡ Limits the ability of the network to provide competing entertainment and other data-rich services
- But caps are easy to ***sell***: Most of the public doesn't realize just how damaging they may be

Diffusing this Conflict: Fairness

- The first step: limit the ***damage*** heavy users can do on normal users
  - ➡ Heavy P2P users can affect light users by interfering with their traffic
  - ➡ Its the result of a mismatch:
    TCP is "flow fair" over short durations
    P2P is long-duration and several flows:
    One P2P user can significantly outcompete many interactive users
- ***Long duration fairness*** solves this problem:
  - ➡ Allows the P2P users to fight amongst themselves, but not affect normal users
  - ➡ Does not affect the direct cost of the P2P bandwidth, but eliminates many of the externalities of P2P users
- Comcast has switched to a very clever QOS based fairness method
  - ➡ When there is no congestion, there is no management
  - ➡ When a network is within 70% of the limit:
    All heavy users (~50% of rated bandwidth over ~15 minutes) are placed into a lower quality of service category
  - ➡ Now under actual congestion:
    - ➡ light users experience ***no*** effects
    - ➡ Heavy users still receive service as long as the light users alone don't occupy all the bandwidth

Diffusing this Conflict: Edge Caches

- P2P systems can work with ***edge caches:***
  - ➡ Caches located ***out of path*** in the minimum-cost position of the ISP's network
    - ➡ They act just like any other node in the P2P system, except they are long lived, high bandwidth, and preferentially serve the ISP's customers
- Unlike HTTP caches, edge caches have deployability advantages:
  - ➡ If they fail, there is no impact:
    - ➡ Unreliable means they can be ***cheap***: "Disk is cheap, ***storage*** is expensive"
  - ➡ They are ***partially deployable***:
    - ➡ ISPs benefit from deployment, but the system still works when they don't exist
- Acts to ***both*** minimize externalities ***and*** save costs
  - ➡ Now only one copy across the ISP's access network, rather than ***N*** copies
  - ➡ ISPs see the benefits of caching, content providers see the benefits of P2P, and the customers get their movies...

So Why Focus on Edge Caches?

- Emphasized with all sides on the conflict:
  - ➡ Content providers want to minimize cost, and already see a pathway with P2P
    - ➡ As a user, I like Netflix's streaming service but the cost does add up..
  - ➡ ISPs have significant network management problems
    - ➡ At ICSI, we've seen what 2 unauthorized P2P users can do to our bandwidth usage
- Look for solutions which would benefit both sides
  - ➡ Conflicts don't have to end with both sides losing: The best outcome is if both sides can ***win***
  - ➡ Which is why I've been focusing some of my effort on what sort of system could be deployed to benefit all parties

The Malcode Wars...

- One of the most pernicious threats to computers is ***malicious code***:
  - ➡ Programs written to automatically compromise victim systems to further an attacker's objectives
- Many different forms:
  - ➡ Virus: A self propagating program that infects files
  - ➡ Worm: A self propagating program that spreads through the network
  - ➡ Trojan: A program which exploits the user's system but does not spread
  - ➡ Botnet: A program/communications system that allows an attacker to easily control hundreds or thousands of compromised systems

But Malcode ***is not the problem...***

- The problem is the ***usage*** of malicious code:
  - ➡ In the 80s and 90s, most malcode was simply for amusement:
    it didn't necessarily do damage or cost money
  - ➡ Malcode is just ***technology***...
- But recent malcode has arisen from two applications:
  - ➡ The ***for-profit*** botherders:
    Malcode is now a profitable business
  - ➡ The ***for-espionage*** botherders:
    Malcode as a targeted weapon

My Viewpoint

- This is an area of primary research for me
  - ➡ Involvement in worm/malcode defense since 2001
- Much of my interest in decision cycles arises from this work
  - ➡ Worms operate on non-human timescales: which requires automated defenses
    - ➡ Communication and the speed of light matter: Some worm attacks ***can not*** be blocked with collaborative defenses

Understanding the Problem

- Malcode for Profit:
  - ➡ Driven by economic factors
    - ➡ "Open" economies vs vertically integrated institutions
  - ➡ Disruptible? by economic and technical factors
    - ➡ Economic infiltration:
      The ***Dark Market*** Takedown
    - ➡ Technical infiltration:
      Botnet infiltration
- Malcode for Espionage:
  - ➡ Far less visibility, but some:
    The ***ghostnet*** experience

The Criminal Economic Underground

- Criminals are in business to ***make money***:
  - ➡ "Why do you rob banks?"
    "Because that's where the money is."
    -falsely attributed to Willie Sutton
- But, like all activity, criminals can benefit from ***specialization***
  - ➡ Don't do everything poorly, do one thing ***well***
- But Specialization requires ***Organization***:
  A way of uniting differing talents

Open-World Economies

- "Open World" economies:
  Economic systems with limited/no restrictions on participation
  - ➡ If you join, you're in
- Advantages:
  - ➡ Allows the maximum benefits of specialization and entrepreneurship
    - ➡ Benefits are greatest when they are largest and most open
- Disadvantages:
  - ➡ Easier for the feds to infiltrate
  - ➡ Easier for parasites ("Rippers", others who prey on their fellow thieves…)

Vertical Integrated Economies

- A single affiliated group which integrates the different expertise
- Advantages:
  - ➡ Harder to infiltrate
  - ➡ Lowers the costs of doing business
    - ➡ In house full time experts are generally cheaper than external “consultants”, ***if*** you can keep your experts efficiently occupied
- Disadvantages:
  - ➡ Limits the available expertise

Reputation in Open Economies

- If you want to be a good criminal, you need a reputation for honesty
  - ➡ Just as in eBay, large economies with infrequent pairwise transactions require reputation systems
- Reputation is developed as a property in a community
  - ➡ If identities are easy to create, only ***positive*** reputation systems work
    - ➡ Only deal with known “good guys”
  - ➡ If identities are hard to establish, ***negative*** reputation also applies
    - ➡ Don’t deal with known “bad guys”
- Open economic systems live or die on their reputation management

DarkMarket.ws

- DarkMarket.ws was one of the “open world” cybercriminal marketplaces
  - ➡ Provide not just forums, but reputation services, marketplace activities, etc
    - ➡ “Your one stop shop for cybercrime” might well be the slogan of these sites
- One of the people with operator status was “Master Splynter”
  - ➡ Master Splinter is the sewer-rat sensei of the Teenage Mutant Ninja Turtles…
- Over a 2+ year undercover investigation, Master Splynter rose from being a random “spammer” to one of the administrators in charge of the site
  - ➡ One of only four major english language marketplaces at the time

Competing Forums and “Max Vision”

- The problem with large criminal marketplaces is ***mistrust, uncertanty, and menace***:
  - ➡ You can have “Rippers”: criminals who pray on other criminals
  - ➡ You can have informants or infiltrators
- Max Ray Butler, aka Max Vision, aka Iceman, had the vision to provide a single unified marketplace as a one-shop stop to address these issues: ***CardersMarket***
  - ➡ He hacked the ***DarkMarket*** server and believed that Master Splyntr was a fed based on logged IP…
  - ➡ But he failed to convince others of this…
    - ➡ Mostly because there is a lot of mistrust of Max Vision:
      His method was to hack other carder sites, suck down the database, move all the users/import all the data into CardersMarket, then wipe the data from the original site
    - ➡ As a result, others claimed Max Butler was a cop!

Both Takedowns Were Successful

- Max Butler was arrested and convicted
  - ➡ Causing CardersMarket to go dark, which had previously attacked the other marketplaces
- There was a flurry of arrests of DarkMarket.ws users
  - ➡ Master Splyntr posted a “F-the-feds, I’m closing down” message when other admins were arrested
  - ➡ But Master Splynter was really J. Keith Mularski, an FBI agent…
    - ➡ Who says feds don’t have a sense of humor? His name was ***advertising*** that he was a rat!
- Put a huge bite out of the English-speaking hacker economy

Takedowns and the Open World Economy

- These takedowns work very well against these open-economies ***for some players***:
  - ➡ Feeds mistrust and menace into the system
  - ➡ The feds have done this multiple times with multiple marketplaces
- The problem is ***regulatory arbitrage***:
  - ➡ Don't be a cyber criminal in the United States
    - ➡ Unless you are under 18, of course. If you are under 18, just don't hack federal and state computer systems...
  - ➡ Don't be a cyber criminal in Turkey if you piss off the authorities
  - ➡ Don't worry about it if you are in Russia, Ukraine, India, etc…
  - ➡ A general limit of all law enforcement strategies when the crimes are location-agnostic

Open Question: Disrupting the Minnows...

- Not ***all*** of the criminal ecology is location-agnostic
  - ➡ The process of "cashing out" often requires location-based actors, commonly low level mules
    - ➡ Bank account owners as "business representatives"
    - ➡ Trans-shippers/repackagers who receive stolen goods
      - ➡ Amazon won't ship a laptop to the Ukraine, making this an exciting work-at-home opportunity
    - ➡ Someone has to ship the Viagra
  - ➡ Unfortunately, these tend to be low skill roles
- Open questions:
  How much can be understood about this portion of the criminal ecology?
  Is it possible to scare ***enough*** potential minnows to make a difference? Inject ***menace*** into that ecology?

Extending ***mistrust and uncertanty*** to the Bots Themselves...

- ~~Un~~fortunately, TPM for PCs failed
  - ➡ As a result, there is no way to ***attest*** that the PC really is what it says it is: you can't validate that the code running on the remote system is only the code the author intends
- Good for Botherders:
  - ➡ TPM was supposed to prevent malicious code (ha)…
- But actually really bad for Botherders:
  - ➡ There is ***no way*** for the botherder to have assurance that their bot is running in a pristine environment
- This enables Botnet ***infiltration***:
  - ➡ The good guys run the bots within contained systems to see what they do: Allows the good guys to directly assess the bad guy's objectives, such as finding out exactly what spam is sent
  - ➡ For a non-secured botnet, the good guys have even been able to ***modify*** the botnet traffic

End States and Botnet Infiltration

- Detection and counter-detection of botnet infiltration has defender-favorable end states
- Detect good-guy introduced monitoring code:
  - ➡ Requires the bad guys to solve the AV problem...
- VM detection:
  - ➡ Code can ***always*** detect that it is running in a VM…
    But who says our captured bots have to run in a VM?
- Human detection:
  - ➡ Is there a human on this computer (with a variety of heuristics)?
    We can always put humans on the computer if we only need one or two bots in a particular botnet
- Network behavior detection:
  - ➡ The one open arms-race: detecting whether the containment mechanism exists
- This seems an arms race ***worth fighting***: the good guys have inherent advantages

Malcode for Espionage

- Once you get a bot running on a victim, you can do whatever you want…
  - ➡ So why just send spam?
    The technology can just as well send secrets
- We have seen the rise of targeted "Malcode for Espionage"
  - ➡ The basic MO:
    Send an email to your target.
    This email contains an attachment that includes an exploit for the program
    The exploit infects the victim when the attachment is viewed
    - ➡ "Spear Phishing" with malcode
    - ➡ The bad guys spend a huge amount of work on "usability" here:
      The email is forged to come from a known person of the victim
      The attachments are often specially crafted to be of interest to the victim
      The attachments are tested against AV programs to avoid detection

The ***GhostNet*** Incident

- The office of the Dalai Lama was under attack from targeted malcode
  - ➡ In a rare break, they brought in outsiders from Toronto, InfoWar Monitor, F-Secure and Cambridge to investigate
    - ➡ And the researchers were allowed to ***publish*** their findings
- The bots were controlled through a central web site…
  - ➡ With ***guessable*** password protection!
- The good guys got onto the web site
  - ➡ And saw what there was to see:
    NGOs, embassies, and lots of other high profile, chinese relevant espionage targets
- Targeted attacks are being ***actively used*** for espionage

This is a ***Tough*** Problem

- Attackers are using ***novel*** malcode:
  - ➡ Stopping “Known Bad” doesn’t work anymore
    - ➡ Billions of dollars of antivirus software is basically obsolete
- Attackers are being ***well targeted***:
  - ➡ They know how to target specific individuals
  - ➡ A lot of work on bad-guy usability
- This looks to be a very interesting arms race going forward...

Security in my Everyday Life: My Financial Protocols

- I've developed a complex set of financial protocols for my everyday life
  - ➡ How I handle my wallet, my credit cards, my debit cards, etc…
  - ➡ A detailed security analysis went into constructing these protocols
- Designed around a mostly opportunistic adversary:
  - ➡ Criminals who aren't targeting me, but rather targeting everybody or anybody
- And designed around minimizing ***my*** costs in the case of a breach:
  - ➡ I don't want perfect security, I just want to be able to stick someone ***else*** with the bill
- Focused on ***what*** to protect
  - ➡ Which dictates the ***how***

Part 1: Passwords

- Passwords are a total pain:
  - ➡ They are hard to remember
  - ➡ They are easy to steal
- I try to avoid them when possible:
  - ➡ SSH public key authentication everywhere:
    Also resists some damage from host compromise
    - ➡ Making public key systems ***as usable as possible*** increases convenience and security
- Otherwise, I write them down ***in my wallet***
  - ➡ If someone has access to my wallet, they have access to my computer ***anyway***

Part 2:
Credit Cards

- I'm actually very cavalier with my credit cards:
  - ➡ I use them for almost all my purchases
  - ➡ I will use them online, and, if necessary, in email
  - ➡ I take no special care to protect them
- Its not because of the threat of compromise, but because of the ***damage***:
  - ➡ Initially, it is the ***credit card company's*** money that is jeopardized
    - ➡ Why credit card companies have good fraud protection
  - ➡ In the end, it is usually the merchant who allows a bad transaction that is responsible
  - ➡ And I have 2 credit cards, so if one dies, the other is still good...

Part 3: Debit Cards

- My ATM card is treated ***very*** differently:
  - ➡ It does ***not*** work through the credit card system: it is “ATM only”, not a “Visa/Mastercard Debit Card”
    - ➡ I had to request this special from my bank
  - ➡ I ***only*** use it at physical ATMs which are built into banks or similar
    - ➡ I don’t want it to get caught in retailer data breaches
  - ➡ I physically inspect the ATM for skimmers before using
- Why so paranoid?
  - ➡ Until the disputed transaction is settled, it is ***my*** money thats on the line, not the bank’s
    - ➡ Eventually the bank would probably have to make things right, but in the mean time…
  - ➡ Driven from an ***economic*** analysis of the cost of a breach and a ***tactical*** analysis of the attackers’ opportunities to achieve the desired data
- Focused on limiting exposure ***regardless*** of attacker tactics

Part 4:
~~Online~~ Banking

- I don't do online banking
  - ➡ I pay a couple bucks a month by sending checks in the mail (dropped in a USPS mailbox):
    - ➡ Attacking checks in the mail is O(N) and proximity-limited and high risk
    - ➡ Attacking electronic banking through malcode on users' computers is O(1) and proximity agnostic
      - ➡ Which would attackers prefer to do?
  - ➡ Or sometimes pay by phone/kiosk with a credit card
    - ➡ Get the 1% kickback from the credit card company as a bonus...
- I do very ***very*** limited access to my brokerage account
  - ➡ Once every few months
    - ➡ Active trading is a great way to lose a ***lot of money quickly***
  - ➡ Ideally, I boot from an Ubuntu live CD into a trusted-boot environment
- I may try to keep malcode off my computer
  - ➡ I run a mac, I use Opera as my primary browser, I keep things up to date… But when it is ***my*** money at stake, I have a very paranoid attitude

Why Wall Street Can't Work

- The Wall Street Firm's Stated Strategic Goals:
  - ➡ The firm should make real, long-term profit
- In order to ***incentive*** the employees who are motivated by money, employees are compensated based on performance
  - ➡ However, effectively all compensation is based on the ***appearance*** of ***short-term*** profit
    - ➡ Can't compensate based on actual, long term profit on a timescale less than measured:
      Wall-Street employees does not want to be paid a bonus ten years later
- Competition ***between*** firms is often predicated on the ***appearance of short term profit***
  - ➡ Worse: many firms have ***innate mistrust*** in their corporate DNA:
    - ➡ "You Eat what you Kill": Your competition is in the next office over
- Thus the Employee's Strategic Goal: ***Appear*** to be profitable in the ***short run***

My perspective

- Out of necessity
  - ➡ I have ***enough*** savings that I need to worry about investments, asset allocation, and all that good stuff
    - ➡ And have gotten lightly burned in the process
- Out of curiosity
  - ➡ I have a ***casual*** interest in Economics
    - ➡ Money impacts security so much that its a good field to study
  - ➡ Anything involving ***this much money*** must be crooked in interesting ways
- Out of having a ***backup plan***:
  - ➡ There are always roles for “Rocket Scientists” in Wall Street:
    If the NSF stops funding my research, I should understand how to fleece Manhattan
- Mostly from the informal viewpoint of my evil twin,
  but with the Boyd viewpoint for understanding specific institutional abberations

Start with the Incentives and Strategy:

- As an ***individual*** in the institution, your strategic goal is to make ***more money***
  - ➡ Its often not the money itself, its that ***money is a way of keeping score***
- Thus the goal of a ***rational*** individual is to create ***the appearance of profit***
  - ➡ "Jesus is coming, everybody look buisy"
  - ➡ The goal of a ***slightly less ~~(cold, calucalting bastard)~~ enlightened*** individual is to do things which the individual ***believes*** will create a profit

Consider the ***Resources and Opponents***

- Smart people, lots of computers, and marketing departments…
  - ➡ All players have all three
- Many competitors have ***deliberately disfunctional*** structures
  - ➡ ***Mistrust*** rather than ***harmony*** in internal dealings
    - ➡ How can you run a company when your employees compete with each other?
  - ➡ ***Uncertanty*** is countered with ***"Certanty"*** in the form of models
    - ➡ Unfortunately, these models aren't actually adaptable and agile
    - ➡ But this is accepted practice by the big pools of money...
- And big pools of money out there…
  - ➡ The goal is to attract the giant pools of money to your institution…
  - ➡ More money -> More return -> more money for you!

Develop the Tactics

- The tactical objective is to ***appear*** profitable
  - ➡ Especially in normal times
  - ➡ It is OK to fail later…
    - ➡ Just say "oh well, black swan, nothing we could do about it…" as you sail away in your yacht
- Extract value continuously
  - ➡ Some fraction of ***each year's profits*** + management fees
    - ➡ Stick that in something with more long term value: Treasuries, yachts, Florida swampland...
- But the question is, how to appear profitable?
  - ➡ While still being ***legal***:
    You don't want to spend the next 150 years in PMITA Federal Prison with Bernie Madoff as your cellmate

How To
Appear Profitable

- Two tactics to ***appear*** profitable:
  - ➡ Hide the Long Tail Risk
    - ➡ The Lake Woebegone Syndrome...
  - ➡ "Profitable" Zero-Sum Games
    - ➡ The Quantum Physics of Options and Derivatives
- One strategy to ***ensure deception***:
  - ➡ Regulatory and Ratings Arbitrage
    - ➡ "But Mommy, Daddy said I could…"

...Where All Funds are Above Average

- If you are running a hedge fund or active mutual fund, your performance is evaluated against an index/ benchmark and your peers
- Yet a puzzle:
  If all else is equal, it is very ***very*** hard to outperform the index…
  - ➡ There is a ***lot*** of math behind why index funds work
- So why are all the hedge funds above average?
  - ➡ Well, there's the survivor bias:
    - ➡ Only the funds that do well get benchmarked
  - ➡ And the suckers:
    - ➡ Enough suckers in a market can allow the "smart" to achieve higher than average return

But Return is only Part Of the Story

- The ***all else equals*** part is the exception:
  - ➡ The returns are only part of the story…
    The index is the optimum long term return ***at a given risk and liquidity point***
- You can be ***lucky***:
  - ➡ After all, you do have a 50% chance of being above average
- You can beat the market if you sacrifice ***liquidity***:
  - ➡ This is one of the primary tenants of Berkshire/Hathaway:
    Buy companies and hold them forever…
    - ➡ The other tenant is get lots of leverage by being an insurance company because if you can price risk ***right*** it becomes much cheaper than ordinary borrowing
- You can ***seem*** to beat the market if you increase ***risk***:
  - ➡ Especially “long tail” risk: low probability but highly catastrophic risk
  - ➡ Most of the time, everything is great
    - ➡ But if bad happens, its ***bad***

Capital Decimation Partners

- The "Capital Decimation Partners" hedge-fund strategy:
  - ➡ Short out-of-the-money put options
  - ➡ Until the market collapses, you make a fortune!
    - ➡ But when the market collapses, you get wiped out: it takes on a huge amount of long-tail risk
    - ➡ From "Risk Management for Hedge Funds", Andrew Lo,
- There are many, many other strategies with the basic suicidal property:
  - ➡ 90% of the time, the fund returns an above-market rate (e.g. a 10% boost in annual return)
  - ➡ 10% of the time, the system implodes and the fund's value goes to 0!
    - ➡ Long term guaranteed rate of return of -100%
- You can always run such strategies inadvertently
  - ➡ People are very good in seeing patterns that aren't there and underpricing long-tail risk
- Secrecy blinds your "opponents":
  - ➡ Both other hedge funds and your own investors!
    (They don't have observations/data)
  - ➡ Even with "disclosure" there is too much going on for someone to really be able to make a rational decision

I Experienced This Personally: Schwab's Bond Funds

- Schwab has two short term bond funds:
  - ➡ The promise: High liquidity and very low risk
  - ➡ SWBDX: Schwab Short Term Bond Market fund
    - ➡ >50% AAA rated, less that <5% below investment grade
    - ➡ All investments mature in <2 years
  - ➡ SWYPX: "Yield Plus" Bond Fund
    - ➡ <25% below investment grade, but still B rated
  - ➡ The same managers ran both funds, and they were sold as an alternative to money market funds (very liquid, very safe)
- It should be ***almost*** impossible to lose money with this investment strategy…
  - ➡ After all, just hold the bonds to maturity: even ***junk*** bonds will pay out unless the issuer goes bankrupt
  - ➡ Additionally, these are all very ***liquid*** investments

My Lesson Learned…

- In March of 2008, I saw a news report about "trouble with a Schwab bond fund…"
  - ➡ Log in and see that SWBDX lost 5% of value:
    I pulled my money out immediately
    - ➡ My stock holdings can bounce up and down, but bonds are a different risk profile
  - ➡ Between January 2008 and today, it lost >10%!
    - ➡ How could this happen?
- But the real nightmare, SWYPX
  - ➡ ***Lost ~50%*** of its value since January 2008!
    An over $12B mutual fund completely imploded: Almost everyone has pulled their money out
    Many ***many*** investors lost 25% or more of their investment
- The ***dog bites man*** conclusion: Fund managers will not act in the interest of the fund, but in their own interest
  - ➡ Worse, the lead manager ***is still employed and running (what's left) of the funds!***

The Strange Implication...

- You want to run as large a fund as possible, because the size dictates your profitability
  - ➡ For a mutual fund: .5% fees on $10B is $50M a year
    - ➡ Plus you can make money loaning out securities to short-sellers
  - ➡ For a hedge fund, make that 2% fees and 20% of the profits...
- Anyone who inadvertently or overtly runs suicide strategies in a hedge fund can make it ***appear*** to be highly successful by taking on long-tail risk
  - ➡ And therefore attract more investors
  - ➡ And those who do it overtly are likely to be doing ***a better job of it!***
- At least some of your competitors ***will*** be running suicide strategies
  - ➡ Whether they know it or not
- QED: If you want to be a successful fund manager, ***it is in your interest to drive your fund into the ground!***
  - ➡ Make it legal with a little willful ignorance…
  - ➡ And when it all fails go, "Oh well, black swan, we could have never predicted that…"

Tactic 2: Derivatives as Quantum Zero Sum Games

- The stock and bond markets are net-positive-sum games:
  - ➡ Real companies make real profit which flows into the system
    - ➡ However, the system size is ***finite***: As a wall-street firm, you can't arbitrarily decide that there should be an extra 1 million shares in Microsoft on the market...
- Derivatives, options, credit default swaps, etc... are net zero sum games:
  - ➡ Side bets: "I pay you $X, if condition Y happens, you pay me back $Z"
  - ➡ For every dollar "earned", a dollar must be lost:
    overall the system ***can not make money on derivatives!***
- But derivatives are "Quantum Foam":
  - ➡ In quantum dynamics: particle-antiparticle pairs are allowed to be created and annihilated without violating conservation laws
  - ➡ In finance, derivatives are the same thing:
    Two parties get together and decide to create a derivative contract
    - ➡ Thus, yay, the potential pool is ***infinite***

So Why
So Many Derivatives?

- AliceCorp and BobCorp decide on a derivative trade:
  - ➡ AliceCorp pays BobCorp $1M a quarter for the next 5 years. If MegaCorp defaults on $100M in bonds within the next 5 years, BobCorp will pay AliceCorp $100M
    - ➡ This is a Credit Default Swap
    - ➡ Depending on Bobcorp's credit rating, it may not have to put up any collateral
- So why do this?
  - ➡ The stated reasons:
    - ➡ AliceCorp owns $100M of MegaCorp bonds, and wants to protect itself
    - ➡ BobCorp is writing an insurance policy
  - ➡ But AliceCorp doesn't own MegaCorp bonds, and BobCorp is not acting like an insurance company does...

The likely explanation: Everybody is A Winner

- Alice (@ AliceCorp) and Bob (@ BobCorp) can both say what a good job they are doing...
  - ➡ So they both deserve their huge bonuses...
- Alice's accountant and Bob's accountant can say these are possibly profitable:
  - ➡ The Wall Street solution for ***uncertanty*** has not been observation and adaptability but ***modeling***
  - ➡ The Black/Schols model for option pricing has many fudge-factors
    - ➡ So both AliceCorp and BobCorp can book a profit
    - ➡ No valid ***observations*** -> no valid decision cycle
- But wait, there's more...

Chaining and Counterparties...

- Suppose CarolCorp comes along and will sell BobCorp the same CDO for $900K a quarter
  - ➡ Now BobCorp ***buys*** the CDO from CarolCorp, hedging the CDO ***sold*** to AliceCorp...
    And pocketing $100K every quarter
- Now suppose MegaCorp starts to get in trouble...
  - ➡ Now AliceCorp ***sells*** a CDO to DaveCorp for $1.1M a quarter...
- We now have $300M in “insurance” coverage on a $100M bond!
  - ➡ DaveCorp bought from AliceCorp which bought from BobCorp which bought from CarolCorp
  - ➡ This is why the “notional value” (the amount outstanding) can be so outrageously high for options and derivitives

But the ends ***must*** be suckers

- In a long chain of options, one end or the other (or both) ***must*** be suckers...
  - ➡ Each link in the chain means the price stretched out away from reality
    - ➡ Chains may be worse than open markets: Open markets only create suckers on ***one side*** of the chain, not both...
- But if the sucker at the end breaks, the whole ***chain*** fails
  - ➡ If CarolCorp fails and the chain has to pay off, now ***everybody*** is on the hook but ***nobody*** has the money for it
    - ➡ This is “counterparty risk”
  - ➡ If DaveCorp fails, now AliceCorp is back on the hook for $1M a quarter...
- And you would expect that the insuring sucker would naturally insure a lot...
  - ➡ And it was called ***AIG Financial Products***

This is How AIG Nearly Totaled The Global Economy

- There were some massive amounts (trillions?) of notional value credit default swaps
  - ➡ Many insuring against the default of bonds charitably described as “Toxic Waste”
- And all chains lead to AIG
  - ➡ The effectively unregulated financial products division in London which wrote tens of billions of dollars of CDSs
  - ➡ When things started going south, AIG had to put up collateral
    - ➡ Which wasn’t enough, because they wrote far more “insurance” then they could ever pay out
- When things went bad, the Treasury chickened out...
  - ➡ Rather than letting AIG fail and bailing out the counterparties... The treasury bailed out AIG directly and made good on the insurance
    - ➡ Allowed all the counterparties *cough* Goldman/Sachs *cough* to look like they were geniuses who made lots of money...

Companies vs Regulators

- So how did AIG get away with getting into a position to destroy the global economy?
  - ➡ Simple: The institutional incentive is to ***ensure*** a lack of oversight: the best regulation is no regulation
    - ➡ Thus for ***this*** conflict, the two groups are the financial institutions and the regulators
- For an institution like AIG, each sub-business has its own regulator
  - ➡ Or, is like AIG Financial Products which had no effective regulator
- But there is the overarching regulator for the entire business
  - ➡ Need to select the regulator who best serves the business's objectives

The "*But Mommy Said I Could*" Theory of Regulatory Oversight

- The dirty little secret:
  Federal regulators are funded by the companies they regulate
  - ➡ So although their ***stated goals*** are in ***harmony***, their actions are in ***mistrust***
    - ➡ Whenever your opponents act out of mistrust, they become exploitable
- So companies will play regulators against each other to find the most compliant regulator
  - ➡ A trick every child knows: exploit ***mistrust*** and ***divergent interests*** between your parents

A Long and Glorious Tradition...

- The Savings and Loan Crisis:
  - ➡ The S&L regulator was significantly weaker than the bank regulator.
  - ➡ After failing, the S&L's regulator was dissolved/restructured as the "Office of Thrift Supervision" (OTS)
- AIG
  - ➡ Ran a small thrift in order to be regulated by the OTS, which, once again, became regarded as a weak regulator
- Another weak regulator was the Office of the Comptroller of the Currency
- As long as there are competing regulators, this process will continue

So What Does This All Mean For Me?

- Index funds and treasuries are the only things I can trust
  - ➡ I can't beat the index, so why try?
  - ➡ I am already "long on the federal government" by being a US Citizen...
- Any financial advisor who does ***not*** advise index funds, treasuries, etc... is acting in ***his*** interest, not mine
  - ➡ Corollary: The only major financial advisors I'm willing can trust are in the same boat:
    One of the primary successes of Berkshire/Hathaway is that Warren Buffett and Charlie Munger are ***long term*** shareholders

And for society?

- Unless compensation is fundamentally changed, these problems will recur again, and again, and again...
  - ➡ People act in their own ***self interest***
  - ➡ And ***any*** compensation scheme based on short term perceived profits ***will*** be corruptable
- Only ***long term*** compensation may work:
  - ➡ E.G. The bonus is in the form of stock where only 5% a year may be sold
    - ➡ Acts to ***realign*** the employee's incentives to those of the long-term shareholder/owner
  - ➡ Why Berkshire/Hathaway works… Its structured around long term ownership
- Otherwise? Forgettaboutit...
  - ➡ More S&Ls. More LTCMs. More AIGs. Same song, different day

Conclusions...

- 'Hi, I'm Nick, and I'm a really suspicious bastard.'
- Security is about how ***people*** behave
  - ➡ Their motivations and intentions
  - ➡ The nature of their organizations
- Understand and attack your opponent's decision making processes
- Security is ***fun***
- I hope everyone found this interesting...