

# **“The More You Spam The Less I Care:” Psychological Decision Making Heuristics and Their Impact On Anti-Spam Activity**

Joe St Sauver, Ph.D. (joe@uoregon.edu)

Senior Technical Advisor

Messaging Anti-Abuse Working Group

MAAWG 14th General Meeting, Ft Lauderdale, Florida

September 22nd-24th, 2008

<http://www.uoregon.edu/~joe/maawg14/>

**Disclaimer:** All opinions expressed are solely those of the author and do not necessarily represent the opinions of any other entity or person.

# **I. Introduction**

# **The Inspiration for This Presentation: Professor Paul Slovic's Talk on Genocide**

- For those who haven't met Paul Slovic and who may not be familiar with his work, Paul's a psychology professor at the University of Oregon; President of Decision Research Institute, Inc.; and an acknowledged expert on the perception of risk and the psychological heuristics of decision making.
- On the Friday before Memorial Day 2008, Paul delivered a talk at the University of Oregon which I attended. It was entitled "If I Look At the Mass I Will Never Act: Psychic Numbing and Genocide," following up on a week-long seminar on genocide held by the Auschwitz Institute for Peace and Reconciliation.
- His talk highlighted a paradoxical phenomena, namely that while most people will go to great lengths to help an individual victim, they may be completely indifferent to the plight of a group whose members may be in equally dire circumstances.

# **No, I'm NOT Comparing Spam to Genocide, and No, Paul Slovic Is NOT Responsible for This Talk**

- Let me be very clear right up front: in no way do I mean to imply that spam is as serious a problem as genocide, because it isn't. We may, however, share some common decision making problems when we think about both issues.
- Let me also be very clear that Paul Slovic is not responsible for the content of this talk, I am. Paul's a very careful and methodical person who will usually have multiple studies to support any assertion he advances; this talk is obviously far less formal and far more conceptual, and is merely meant to outline some areas for potential in depth exploration and verification, it is not meant to report definitive results for work already completed.

# Two Scenarios

- *Scenario 1:* An unattended toddler falls into deep water. What would you do? Answer: you'd probably do whatever you possibly could to rescue that child before he or she drowned.
- *Scenario 2:* Millions of refugees are dying in Darfur, in Sub-Saharan Africa. Money, or even just a little world attention, could help save at least some of those refugees. Yet what do we do? Answer: most of us will do nothing.
- Why the difference in response? Why would we risk our own life saving a drowning child, yet do nothing, not even call our Congressman or Senator, to help save children dying in Darfur?
- It's all about how our brains are programmed to make decisions.

# Factors Affecting The Toddler Rescue?

- The risk the child faces is immediate, acute, certain and serious
- We may be the only chance that child may have; no one else is there to save the child
- Rescuing a child is only a momentary investment of our time
- Rescuing a drowning child is “plausible”/“conceivable” -- we can mentally imagine ourselves successfully accomplishing this urgent but fundamentally simple task
- Having been rescued, there’s an excellent chance that the child will successfully grow up to become a productive adult, so our rescue efforts won’t be “wasted”
- We may receive accolades from others for our efforts; on the other hand, if we fail to try, we may become a target of contempt.
- Deep seated paternal or maternal instincts may be somehow triggered, and we may almost at least for a moment forget that we’re saving someone else’s child and not one of our own.
- Time is critical: we act, we don’t think.

# Some Reasons Why We Don't Save The Refugees

- **The refugees haven't been personalized: we don't know them.**  
Quoting Stalin: "One death is a tragedy; a million, a statistic."
- We're overwhelmed and numb. "Where would we even start?"
- The risk the refugees face is chronic, protracted, and uncertain.  
"Things might get better. The famine might not affect everyone."  
We take comfort in unwarranted optimistic uncertainty.
- Saving a refugee isn't quick; they may need help for years.
- We can tell ourselves that someone else might step forward to help them. Distance provides an excuse for us to shirk our global responsibilities. "Aren't those refugees someone else's problem?"
- Because "we can't possibly save them all" we don't try to save even some (heck, we don't try to save even one). How could we be so "unfair," picking one person to save, recognizing that by doing so, we're effectively "selecting" others to be left to die?
- A highly positive outcome for some might be marginalized by the highly negative outcome of the rest. "Nearly all of them died."<sup>7</sup>

# We've Learned to "Pick Our Shots"

- As a species, we've become quite good at making decisions about selectively employing the limited resources at our disposal.
- We've learned adaptation strategies. For example, if one thing doesn't work, we try something else. Non-productive behaviors that repeatedly fail are eventually extinguished.
- We learn to accept that there are some things which cannot be changed. If it does us no good to rail against swarming bugs and a constant cold driving rain, we learn to "tune those annoyances out," working around them as best we can.
- We tend to focus on the things that are local, because for eons, it's been the local things that can hurt or kill us, and it's the local things that we've been best able to influence or change.
- These deeply encoded approaches continue to guide our behavior and our decision making today, just as they did when we roamed the countryside as part of a tribal band, slaying animals for food.
- Oh wait, we are coming up on deer hunting season, aren't we? :-)

## **II. User Decision Making And Spam**

# ***"What In The World Does All This Have to Do With Spam?"***

- Excellent question. While I was sitting in that seminar room listening to Professor Slovic, I suddenly realized:

**We routinely make spam-related decisions using the same sort of odd/seemingly "irrational" or "inconsistent" thinking processes we apply to other decisions.**

and

**If we don't pay attention to those psychological decision making rubrics, eventually we'll lose the war on spam. People aren't computers, and they don't act like computers -- so we shouldn't treat them like computers!**

## In Particular...

- We do really poorly when it comes to reasoning about large numbers, or dealing with “lots” of anything.
- One of the reasons we have trouble with genocide is because we can't wrap our heads around the thought of hundreds of thousands (or millions of people) dying. The sheer thought simply makes us numb, and we begin to avoid thinking about it.
- Similarly, we have trouble with spam because it is hard for us to conceive of hundreds of millions or hundreds of billions of spam per day, or even our trivial “share” of that huge daily spewage. Like genocide, spam can be so overwhelming it makes us numb.
- Yet users still need to make daily decisions about email and spam.

## The Decisions That Some Users Make May Not Seem To “Make Sense”

- **I just got spammed. Should I complain to someone about it?**

*When I only get a few spam, I complain about each and every one of them. But if my inbox is completely swamped with spam, I just delete 'em all. I don't have time to complain about them all.*

- **The unsolicited commercial email I just got is actually advertising something I think I want to buy...**

*No one will know if I buy something I saw in a spam message. And that cute handbag was such a bargain! What could it hurt?*

- **Some of the unwanted messages I receive have an "opt-out" link. I'm *really* sick of getting spam. Maybe I'll just try to "unsubscribe" ...**

*Why would spammers ignore my request to unsubscribe? I'll just try it on a couple of the spam I get, and see if it helps.*

# Many Spam-Related Decisions Actually Are Rational, If Thought About In The Right Way

- Let's begin by considering the title hypothesis of this talk:

**The more you spam me, the less I care.**

- Most of us would probably assume that the exact opposite would be true. Because spam is irritating, the more you spam me, the more irritating I'd find it that experience, and the more irritated I became, the more inclined I'd be to try to "do something" about the spam I receive. Right? *Right?* Maybe not.
- What are a user's options, having gotten spammed?

# Users Have A Lot of Choices

- 1) I could **complain** about the spam I received. This might mean just pushing a “this is spam” button (or I could try reporting spam manually or I could use a reporting service such as SpamCop)
- 2) Alternatively, I might attempt to “**unsubscribe**” from the spammer’s mailing list (although this often “backfires”)
- 3) I might try to “**hide**” from spam by changing my email address
- 4) I might simply **delete the spam I received without opening it**
- 5) I might try augmenting or tweaking my **technical spam filtering**
- 6) I might **use email less often** (or replace it with instant messaging)
- 7) Some users might even be tempted to **open and read their spam**
- 8) Some users may even respond to the spammer's call to action by **buying the product or service that’s being spamvertised**
- 9) Some might even be lured over to the dark side, having been tempted into **becoming a spammer** themselves

# Let's Assume I Complain About My Spam

- **<user clicking on the “this is spam” button>**  
Take ***THAT***, damn spammers! Having done my part to eliminate spam, I logout and go watch football game filled with a self-satisfied glow, confident that my spam is a thing of the past.
- *Time: next day.* I open my inbox, expecting to see it more or less spam free, only to see as much (or more!) spam as the day before. What the ?!@#?!? Why hasn't all the spam “gone away” now that I've complained about it? Maybe it just takes a while to “work”...
- It might be a matter of days, or it might take weeks or months, but I predict that **eventually** an apparent lack of effect will result in many people becoming disillusioned and deciding that there's “no point to complaining” because complaining “doesn't help.”
- Recall what happens to problem solving strategies that don't work: our primitive brains eventually notice, and discard those that don't seem to work. (Users can't spend their whole day doing nothing but complaining about spam!)

# Some Other “Reasons” Why Users Don’t Complain

- Besides “no point to it/doesn’t accomplish anything”...
- Other people are complaining, so I don’t have to do so, too
- If I get a lot of spam, it takes too long to complain
- Afraid of making a mistake (reporting a legitimate message, etc.)
- Have to open the spam to complain; worry about viruses
- Have to open the spam to complain; that tells’ the spammer that I “might be interested” and gives him/her “credit” for reaching me
- Have to open the spam to complain; content is morally disgusting (e.g., explicit adult content) or illegal (child porn)
- Worry that I may be able to be identified as having complained, and the spammer may retaliate against me in particular
- They *said* that I signed up for it; I must have “just forgotten,” or maybe someone else (e.g., a friend of mine) signed me up. If so, it wouldn’t be “fair” or “right” for me to complain. <cough>

# “Congratulations! Complaints Are Way Down!”

- We now know this may **not** be good.
- Are complaints down because:
  - Spam levels have decreased?
  - Spam levels are constant, but our filtering has gotten better?
  - Spam levels are constant and our filtering hasn't improved, but our users are “burned out” and just aren't complaining to us about the spam they're getting anymore?
  - Users are completely ignoring our spam ridden email accounts and are using something else for their email instead?
  - Spammers are somehow identifying “complainers,” and are “list washing” those addresses (and only those addresses) so that complaint rates are down, but most folks are still getting hammered by spam just as hard as ever?
- Make sure you interpret the statistical phenomena you see correctly!

# Potential Action Item/Recommendation

- Assumption: properly formatted and timely spam complaints are helpful, and a useful intelligence source you don't want to lose.
- Check: what steps are you taking to make sure that...
  - it is as easy/painless as possible to file spam complaints?
  - the complaints that users file have the content you need (e.g., timely complaints, full headers, usable attachments, etc.)
  - users know that their complaints *ARE* having an actual impact?
  - users are aware that you appreciate and value the the time they've taken take to file spam complaints (this does not mean sending them boilerplate auto-acks, by the way)
  - users know that spam control *IS* genuinely improving?
- **Recommendation: communicate with your users about the war on spam. Help them technically to complain effectively, make sure you explain that you appreciate their efforts to help, and give them feedback about how the war is going.** 18

# Each Other Potential User Response Should Receive Similar Analysis and Reinforcement

- For example, at least some desperate users may try "unsubscribe" even if they didn't sign up to get a spammer's spewage in the first place. When users attempt to unsubscribe that way, they commonly “just get more spam,” so users quickly learn that attempting to "unsubscribe" can “just make things worse.”
- However, at least in the case of legitimate senders who are actually sending mail that the user requested, users *should* unsubscribe when they no longer want to receive those messages
- As geeks, we might have an excellent sense of which senders we can trust and which ones we can't, but can we/do we convey that information to users in a simple way that will fit into their decision making heuristics? If not, they'll likely employ a “simplified” rule: if you get mail you don't currently want, click “this is spam” even if you asked for that mail at one point in time.

# Some Additional Problems With “Opting Out”

- *Besides* potentially increasing the amount of spam a user gets...
- Trying to “unsubscribe” tends to legitimize the opt-out paradigm forced on us by the CAN-SPAM act
- It flags the opting-out party as being unusually gullible/naive
- Playing the “opt-out game” doesn’t scale/work in a world where mailing lists can be transferred, traded or sold to 3rd parties
- Opt out email addresses can be used to Joe-Job innocent people unrelated to a given spam campaign (such as anti-spammers)
- Web based opt-out can connect an IP address with an email address, thereby increasing the value of the user’s “record” and making it harder to disprove that the user didn’t actually “opt in”
- “Opt out” links might actually drop malware on the user’s system, or display ads on an opt out page might earn the spammer revenue as part of the opt-out process

# Another Old And Well Respected (But Ineffective) Decision Strategy: Try Hiding From The Bad Guys

- Let's assume that our user (you know, the one who unwisely tried to "unsubscribe" from multiple spammer mailing lists), is now receiving a veritable torrent of spam each day.
- Desperate for relief, the user decides on a new strategy: they'll try changing their email address, hoping/relying on the bad guys not discovering and following them to their new address.
- Unfortunately their analysis fails to account for a host of issues that inherently undercut the effectiveness of that strategy, such as:
  - the user may pick an easily guessed/recycled username
  - they may still post their complete email address on web pages where they can be re-scraped (and munging doesn't work)
  - all usernames on the system they're using may be getting periodically harvested and sold by an untrustworthy insider, etc

# So WHY *Do* Users Try Ineffective Strategies?

- Well, maybe you haven't clearly told your users that "trying to hide" doesn't work, nor explained why it doesn't work.
- Desperate for a solution, and not having gotten authoritative leadership from a trustworthy source (e.g., you), they then seek advice from their friends and acquaintances, proceeding to act based on that anecdotal (and sometimes bad) guidance.
- And oh yes: changing usernames is typically inexpensive/free, so often there's no real economic disincentive to surmount.
- **Recommendations: Show leadership. Make sure you talk to your users about things like username changes and why that sort of approach often isn't effective at stopping spam.**
- **Make sure your fee structures send a consistent message. For example, if you believe username changes are a hassle and are not an effective anti-spam strategy, make sure username change fees are consistent with that perspective, albeit with slack for name changes due to marriage or court orders.**

# Users Becoming Resigned

- Beyond trying complaining, unsubscribing and hiding, some users just become resigned to spam, and begin to treat it as an inescapable part of their daily existence.
- This phase is normally characterized by users “just hitting delete.”
- No point to complaining, they’ve already asked the spammer to leave them alone (to no effect), and they’ve tried hiding, so from their point of view, there’s not much in the way of options left, except to simply delete what spam continues to make it through.  
<delete><delete><delete><delete><delete><delete><delete>  
<delete><delete><delete><delete><delete><delete> ...
- From the user’s point of view, this chore, like shoveling snow from the sidewalk all winter if you live up north, is just part of life, a burden, but something which we must accept.
- Repetitive manual process inevitably seems to lead users to explore technological alternatives -- surely there must be a snow blower I can buy that will beat shoveling every day?

# More Recommendations At This Checkpoint

- Are you technically able to detect users who are manually deleting lots of unread messages? Can you tell if those messages all appear to be spammy?
- If so, are you communicating with users in that situation about options that may help them **and** all your other users, like reporting unwanted mail as spam? (“If you get spam, don’t just hit delete, please take a minute to REPORT IT AS SPAM!”)
- Do you have effective options for users who are “shopping for snow blowers” as an alternative to just hitting delete? Again, assume that users will look to you for advice if you’re prepared to offer it! You know (or should know!) what additional filtering software will best augment your production server-side anti-spam filtering.
- At this point you should also have clear ideas about how you want to channel user interest in alternative communication channels (such as instant messaging).

# The Spam Equivalent of the Stockholm Syndrome?

- Having exhausted all options for managing spam, the user may find herself tempted to look at what the spammer is sending.
- Rationalization may be taking place, and the user may be at risk of succumbing to the spam equivalent of the “Stockholm Syndrome:” they may be beginning to identify with the spammer.
- Such a user might begin to actually read the spam that’s still getting through. After all, if the spammer can defeat all these antispam measures, the spammer *must* be pretty smart, right? Maybe I should check out what such a “smart person” is offering. [No, actually you shouldn’t, but some users may be easily influenced and subject to being readily lead astray]
- Having begun to read the spammer’s advertising, the user is then faced with a conflict: what should they believe? Surely the spammer wouldn’t be allowed to continually advertise organ enlargement if it didn’t really occur, right? [wrong, the spammers are just repeating an outright lie often enough that some people begin to believe it *must* be true, a well known propaganda trick]

## *“29% Of Internet Users Buy From Spam”*\*

- If nearly 30% of Internet users actually routinely buy things from spammers, we have a very serious problem on our hands: that level of participation, if true, would be a sign that spam may be beginning to be accepted as “legitimate,” the Internet equivalent of advertising on broadcast TV (even though that analogy is fundamentally and fatally flawed).
- On the other hand, we need to view that number very skeptically, being very alert to possible methodological issues associated with this study, including things as fundamental as how “spam” is defined, and the difference between one time mistakes and routine behaviors.

-----

\* [http://www.darkreading.com/document.asp?doc\\_id=161874&f\\_src=darkreading\\_section\\_297](http://www.darkreading.com/document.asp?doc_id=161874&f_src=darkreading_section_297)

# Why Might Users Be Willing to Buy From Spam?

- An illicit product is being offered (pillz, pirated software, pr0n, casinoz, etc.) not readily available from non-spammy sources
- User doesn't understand why they shouldn't (e.g., might be ripped off; might receive substandard product; might be sharing credit card info with cyber criminals; might be potentially funding hostile entities such as terrorists; sale of “minor” drugs might help to underwrite sale of narcotics and other dangerous drugs, etc.)
- The “lemming” myth: “Everyone does it, why shouldn’t I?”
- The *de minimus* myth: “My tiny little order doesn’t really matter”
- “(Seemed like) a really good bargain” (some spammers are unquestionably excellent (if felonious) salesmen)
- “Seemed just like advertising on TV or radio” (but it obviously isn't since the spammer doesn't help to underwrite the user’s Internet connection, now does he?)

# Recruitment of New Spammers

- Worst of all, having “bought in” to the spam ecosystem by reading spam and maybe even buying something that the spammers are selling, the user is then at risk of becoming a spammer herself. After all, spam “must” work -- didn’t she just buy something from spam? And look how hard spam is to stop! Easy money, too, right? “If you can’t beat them, join ‘em,” etc.
- Now of course the new wanna-be-spammer may NOT see how spammers are getting arrested, or how spammers are getting cheated out of income they’ve “earned” by untrustworthy affiliate programs, or how spam may be fueling unquestionably bad things like drug addiction or international terrorism.
- At root, this recruitment of new spammers is fundamentally a psychological phenomena. It is yet another sign that we’re not doing a good job of pressing home the message that spam is NOT profitable, and spam is NOT a crime that’s safe from prosecution.
- Obviously we need to stop this sort of recruitment of new spammers if we’re to win the war on spam.

# It All Comes Back To Working With Users

- If we can recruit a user to be part of the “anti-spam team,” and we reinforce their feeling of making a difference and being appreciated, we may be able to prevent a downhill slide toward users trying ineffective strategies, user resignation, and eventual user co-option by the bad guys.
- But right now, honestly, how often do you really work user complaints? Sure, it is easier to rely on your own spam traps, or spam trap-driven feedback loop data from other providers, but if you ignore user complaints, I think you’re making a big mistake for some of the reasons I’ve just outlined.

### **III. Why Haven't There Been Any New Anti-Spam Laws Post CAN-SPAM?**

# **I Know Few Anti-Spammers Who Are “Happy” With CAN-SPAM, And Yet...**

- I'm not seeing Congress flooded with new anti-spam legislation.
- Why?
- Again, this may largely be a matter of our flawed decision making paradigms, and the extent to which gradually increasing levels of spam have left us collectively mute, numbed by that flood.
- “The more you spam me, the less I care” strikes yet again.

# A Little More About Why There May Be No New Anti-Spam Bills

- **Industry bodies** (including MAAWG!) aren't drafting and lobbying for new legislation, and legislators themselves may not know what needs to be done differently (and they don't want to "look dumb" if they suggest something that's totally impractical)
- **Because constituents aren't complaining about spam, legislators may have a mistaken perception that spam is a "non-issue"** for them compared to things like the upcoming elections; the wars; economic issues; etc.
- Many **legislators do not use the Internet** and hence do not run into spam issues (sounds almost unbelievable, but it's true)
- Many politicians think that **technology is solving/will solve the spam problem** (and much of the anti-spam struggle is "hidden")
- **The media hasn't held Congress accountable** for a lack of effective anti-spam legislation; spam hasn't been made "Congress's problem" to-date

# Some Possible Reasons Why There Are No New Anti-Spam Bills (2)

- Anti-spam legislation at the **state level** was foreclosed when Congress made spam control almost exclusively a federal issue
- There have been occasional **highly publicized trials** of some spammers, which may make some people believe that we've got enough anti-spam laws already (but obviously we don't)
- **Legitimate mailers**, worried about "bleed through" from inexpertly crafted efforts meant to curtail criminal spammers, may exhibit knee jerk opposition to any/all new anti-spam laws out of worry that a new law may accidentally impact them, too
- **Some legislators may believe anti-spam legislation is futile**, with spammers "working around" or ignore new anti-spam laws as fast as they can pass them, so why even bother trying?
- Spammers work internationally, and hence they sometimes acquire a perceived (but unjustified) mantle of "**untouchability**"

# Many of Those Reasons Are Psychological

- We don't have new anti-spam legislation, not because new legislation isn't needed, but because we've been "psyched out."
- Not only are the spammers continuing to hammer us technically (with spam now running 90-95% of all email at some sites\*), they're also beating us "inside our own heads." We're so mentally whipped we're not even TRYING to take them on with new legislation.

-----

\* [http://www.spamhaus.org/effective\\_filtering.html](http://www.spamhaus.org/effective_filtering.html)

# We've Gotten Too Good at Hiding The Problem

- Because technical measures limit the amount of spam that users see, most users have no sense of just how bad things have become. Their lack of awareness is our fault -- we've done too good a job of hiding the true magnitude of the problem.
- It is tempting to suggest that we need a “no filtering” day once a year, so that users can experience the true magnitude of the spam problem, but in reality, we've already exceeded the point where that would be technically possible.
- If all filtering were to be disabled for even a single day, to bring home the point of just how bad spam has become, that load would crush mission critical systems.

# But The Cost of Spam Is Too Huge To Disregard

- We periodically see estimates of the cost of spam -- for example, Ferris Research has quoted the cost of spam to corporate customers at \$140 billion worldwide, and \$42 billion in the US.\*
- That may seem like a laughably big number, but if you were to spread that over the entire population (checking the Census Bureau,\*\* they say we're around 305,210,497 people in the US, and 6,724,925,242 worldwide), that's only:  
 *$140,000,000,000/6,724,925,242/365=\$0.057/\text{person/day globally}$*   
 *$42,000,000,000/305,210,497/365=\$0.377/\text{person/day in the US}$*
- I think that estimate is **way, way too low**. So why isn't anyone noticing billion dollar hits on our economy? Answer: that money's being taken from us in tiny little slices a billion times a day, so we simply don't perceive it. But what a whack against our economy!

-----

\* <http://www.newswiretoday.com/news/32531/>

\*\* <http://www.census.gov/main/www/popclock.html>

## For Comparison, Some Other Recent Costs

- "Hurricane Katrina cost insurers an inflation-adjusted \$43 billion," [http://money.cnn.com/2008/09/13/news/economy/ike\\_effect/](http://money.cnn.com/2008/09/13/news/economy/ike_effect/)
- "The attack on the World Trade Center will cost New York City \$83 billion to \$95 billion," <http://query.nytimes.com/gst/fullpage.html?res=940DE3DF143EF936A3575AC0A9649C8B63>
- "In February 2008, the Congressional Budget Office projected that additional war costs from FY2009 through FY2018 could range from \$440 billion, if troop levels fell to 30,000 by 2010, to \$1.0 trillion, if troop levels fell to 75,000 by about 2013. Under these scenarios, CBO projects that funding for Iraq, Afghanistan and the GWOT could reach from about \$1.1 trillion to about \$1.7 trillion for FY2001-FY2018."

The Cost of Iraq, Afghanistan, and Other Global War on Terror Operations Since 9/11, Updated July 14, 2008, CRS Report RL33110, page 2.

# Another Reason Why Congress Should Be Paying Attention to Spam

- Consider John Robb's 15 Aug 2008 posting "Open Source Warfare: Cyberwar," ( <http://globalguerrillas.typepad.com/globalguerrillas/2008/08/open-source-war.html> ):

*In contrast to failed US efforts, both China and Russia have adopted the OSW [Open Source Warfare] approach to cyberwarfare. How did they do it? Simply:*

*\* Engage, co-opt, and protect cybercriminals. Essentially, use this influence to deter domestic commercial attacks and encourage an external focus. This keeps the skills sharp and the powder dry.*

*\* Seed the movement. Once the decision to launch a cyberattack is made, start it off right. Purchase botnets covertly from criminal networks to launch attacks, feed 'patriotic' blogs to incite attacks and list targets, etc.*

*\* Get out of the way. Don't interfere. Don't prosecute participants. Take notes.*

# **Thanks For The Chance to Speak Today!**

- Are there any questions?