## Schneier on Security

Essays >

## The Psychology of Security (Part 2)

Bruce Schneier
January 18, 2008

## Return to Part 1

## The Availability Heuristic

The "availability heuristic" is very broad, and goes a long way toward explaining how people deal with risk and trade-offs. Basically, the availability heuristic means that people "assess the frequency of a class or the probability of an event by the ease with which instances or occurrences can be brought to mind. ${ }^{. \frac{28}{}}$ In other words, in any decision-making process, easily remembered (available) data are given greater weight than hard-to-remember data.

In general, the availability heuristic is a good mental shortcut. All things being equal, common events are easier to remember than uncommon ones. So it makes sense to use availability to estimate frequency and probability. But like all heuristics, there are areas where the heuristic breaks down and leads to biases. There are reasons other than occurrence that make some things more available. Events that have taken place recently are more available than others. Events that are more emotional are more available than others. Events that are more vivid are more available than others. And so on.

There's nothing new about the availability heuristic and its effects on security. I wrote about it in Beyond Fear, ${ }^{29}$ although not by that name. Sociology professor Barry Glassner devoted most of a book to explaining how it affects our risk perception. ${ }^{30}$ Every book on the psychology of decision making discusses it.

In one simple experiment, ${ }^{31}$ subjects were asked this question:

- In a typical sample of text in the English language, is it more likely that a word starts with the letter K or that K is its third letter (not counting words with less than three letters)?

Nearly 70\% of people said that there were more words that started with K, even though there are nearly twice as many words with K in the third position as there are words that start with K. But since words that start with K are easier to generate in one's mind, people overestimate their relative frequency.

In another, more real-world, experiment, ${ }^{\frac{32}{}}$ subjects were divided into two groups. One group was asked to spend a period of time imagining its college football team doing well during the upcoming season, and the other group was asked to imagine its college football team doing poorly. Then, both groups were asked questions about the team's actual prospects. Of the subjects who had imagined the team doing well, $63 \%$ predicted an excellent season. Of the subjects who had imagined the team doing poorly, only $40 \%$ did so.

The same researcher performed another experiment before the 1976 presidential election. Subjects asked to imagine Carter winning were more likely to predict that he would win, and subjects asked to imagine Ford winning were more likely to believe he would win. This kind of experiment has also been replicated several times, and uniformly demonstrates that considering a particular outcome in one's imagination makes it appear more likely later.

The vividness of memories is another aspect of the availability heuristic that has been studied. People's decisions are more affected by vivid information than by pallid, abstract, or statistical information.

Here's just one of many experiments that demonstrates this. ${ }^{\frac{33}{} \text { In the first part of the }}$ experiment, subjects read about a court case involving drunk driving. The defendant had run a stop sign while driving home from a party and collided with a garbage truck. No blood alcohol test had been done, and there was only circumstantial evidence to go on. The defendant was arguing that he was not drunk.

After reading a description of the case and the defendant, subjects were divided into two groups and given eighteen individual pieces of evidence to read: nine written by the prosecution about why the defendant was guilty, and nine written by the defense about why the defendant was innocent. Subjects in the first group were given prosecution evidence written in a pallid style and defense evidence written in a vivid style, while subjects in the second group were given the reverse.

For example, here is a pallid and vivid version of the same piece of prosecution evidence:

- On his way out the door, Sanders [the defendant] staggers against a serving table, knocking a bowl to the floor.
- On his way out the door, Sanders staggered against a serving table, knocking a bowl of guacamole dip to the floor and splattering guacamole on the white shag carpet.

And here's a pallid and vivid pair for the defense:

- The owner of the garbage truck admitted under cross-examination that his garbage truck is difficult to see at night because it is grey in color.
- The owner of the garbage truck admitted under cross-examination that his garbage truck is difficult to see at night because it is grey in color. The owner said his trucks are grey "because it hides the dirt," and he said, "What do you want, I should paint them pink?"

After all of this, the subjects were asked about the defendant's drunkenness level, his guilt, and what verdict the jury should reach.

The results were interesting. The vivid vs. pallid arguments had no significant effect on the subject's judgment immediately after reading them, but when they were asked again about the case 48 hours later--they were asked to make their judgments as though they "were deciding the case now for the first time"--they were more swayed by the vivid arguments. Subjects who read vivid defense arguments and pallid prosecution arguments were much more likely to judge the defendant innocent, and subjects who read the vivid prosecution arguments and pallid defense arguments were much more likely to judge him guilty.

The moral here is that people will be persuaded more by a vivid, personal story than they will by bland statistics and facts, possibly solely due to the fact that they remember vivid arguments better.

Another experiment ${ }^{34}$ divided subjects into two groups, who then read about a fictional disease called "Hyposcenia-B." Subjects in the first group read about a disease with concrete and easy-to-imagine symptoms: muscle aches, low energy level, and frequent headaches. Subjects in the second group read about a disease with abstract and difficult-to-imagine symptoms: a vague sense of disorientation, a malfunctioning nervous system, and an inflamed liver.

Then each group was divided in half again. Half of each half was the control group: they simply read one of the two descriptions and were asked how likely they were to contract the disease in the future. The other half of each half was the experimental group: they read one of the two descriptions "with an eye toward imagining a three-week period during which they contracted and experienced the symptoms of the disease," and then wrote a detailed description of how they thought they would feel during those three weeks. And then they were asked whether they thought they would contract the disease.

The idea here was to test whether the ease or difficulty of imagining something affected the availability heuristic. The results showed that those in the control group--who read either the easy-to-imagine or difficult-to-imagine symptoms, showed no difference. But those who were asked to imagine the easy-to-imagine symptoms thought they were more likely to contract the disease than the control group, and those who were asked to imagine the difficult-to-imagine symptoms thought they were less likely to contract the disease than the control group. The researchers concluded that imagining an outcome alone is not enough to make it appear more likely; it has to be something easy to imagine. And, in fact, an outcome that is difficult to imagine may actually appear to be less likely.

Additionally, a memory might be particularly vivid precisely because it's extreme, and therefore unlikely to occur. In one experiment, ${ }^{35}$ researchers asked some commuters on a train platform to remember and describe "the worst time you missed your train" and other commuters to remember and describe "any time you missed your train." The incidents
described by both groups were equally awful, demonstrating that the most extreme example of a class of things tends to come to mind when thinking about the class.

More generally, this kind of thing is related to something called "probability neglect": the tendency of people to ignore probabilities in instances where there is a high emotional content. ${ }^{\frac{36}{6}}$ Security risks certainly fall into this category, and our current obsession with terrorism risks at the expense of more common risks is an example.

The availability heuristic also explains hindsight bias. Events that have actually occurred are, almost by definition, easier to imagine than events that have not, so people retroactively overestimate the probability of those events. Think of "Monday morning quarterbacking," exemplified both in sports and in national policy. "He should have seen that coming" becomes easy for someone to believe.

The best way I've seen this all described is by Scott Plous:

> In very general terms: (1) the more available an event is, the more frequent or probable it will seem; (2) the more vivid a piece of information is, the more easily recalled and convincing it will be; and (3) the more salient something is, the more likely it will be to appear causal. ${ }^{37}$

Here's one experiment that demonstrates this bias with respect to salience. ${ }^{\frac{38}{3}}$ Groups of six observers watched a two-man conversation from different vantage points: either seated behind one of the men talking or sitting on the sidelines between the two men talking. Subjects facing one or the other conversants tended to rate that person as more influential in the conversation: setting the tone, determining what kind of information was exchanged, and causing the other person to respond as he did. Subjects on the sidelines tended to rate both conversants as equally influential.

As I said at the beginning of this section, most of the time the availability heuristic is a good mental shortcut. But in modern society, we get a lot of sensory input from the media. That screws up availability, vividness, and salience, and means that heuristics that are based on our senses start to fail. When people were living in primitive tribes, if the idea of getting eaten by a saber-toothed tiger was more available than the idea of getting trampled by a mammoth, it was reasonable to believe that--for the people in the particular place they happened to be living--it was more likely they'd get eaten by a saber-toothed tiger than get trampled by a mammoth. But now that we get our information from television, newspapers, and the Internet, that's not necessarily the case. What we read about, what becomes vivid to us, might be something rare and spectacular. It might be something fictional: a movie or a television show. It might be a marketing message, either commercial or political. And remember, visual media are more vivid than print media. The availability heuristic is less reliable, because the vivid memories we're drawing upon aren't relevant to our real situation. And even worse, people tend not to remember where they heard something-they just remember the content. So even if, at the time they're exposed to a
message, they don't find the source credible, eventually their memory of the source of the information degrades and they're just left with the message itself.

We in the security industry are used to the effects of the availability heuristic. It contributes to the "risk du jour" mentality we so often see in people. It explains why people tend to overestimate rare risks and underestimate common ones. ${ }^{\frac{39}{}}$ It explains why we spend so much effort defending against what the bad guys did last time, and ignore what new things they could do next time. It explains why we're worried about risks that are in the news at the expense of risks that are not, or rare risks that come with personal and emotional stories at the expense of risks that are so common they are only presented in the form of statistics.

It explains most of the entries in Table 1.

## Representativeness

"Representativeness" is a heuristic by which we assume the probability that an example belongs to a particular class is based on how well that example represents the class. On the face of it, this seems like a reasonable heuristic. But it can lead to erroneous results if you're not careful.

The concept is a bit tricky, but here's an experiment that makes this bias crystal clear. ${ }^{40}$ Subjects were given the following description of a woman named Linda:

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in antinuclear demonstrations.

Then the subjects were given a list of eight statements describing her present employment and activities. Most were decoys ("Linda is an elementary school teacher," "Linda is a psychiatric social worker," and so on), but two were critical: number 6 ("Linda is a bank teller," and number 8 ("Linda is a bank teller and is active in the feminist movement"). Half of the subjects were asked to rank the eight outcomes by the similarity of Linda to the typical person described by the statement, while others were asked to rank the eight outcomes by probability.

Of the first group of subjects, $85 \%$ responded that Linda more resembled a stereotypical feminist bank teller more than a bank teller. This makes sense. But of the second group of subjects, $89 \%$ of thought Linda was more likely to be a feminist bank teller than a bank teller. Mathematically, of course, this is ridiculous. It is impossible for the second alternative to be more likely than the first; the second is a subset of the first.

As the researchers explain: "As the amount of detail in a scenario increases, its probability can only decrease steadily, but its representativeness and hence its apparent likelihood may increase. The reliance on representativeness, we believe, is a primary reason for the
unwarranted appeal of detailed scenarios and the illusory sense of insight that such constructions often provide. ${ }^{41}$

Doesn't this sound like how so many people resonate with movie-plot threats--overly specific threat scenarios--at the expense of broader risks?

In another experiment, ${ }^{42}$ two groups of subjects were shown short personality descriptions of several people. The descriptions were designed to be stereotypical for either engineers or lawyers. Here's a sample description of a stereotypical engineer:

> Tom W. is of high intelligence, although lacking in true creativity. He has a need for order and clarity, and for neat and tidy systems in which every detail finds its appropriate place. His writing is rather dull and mechanical, occasionally enlivened by somewhat corny puns and flashes of imagination of the sci-fi type. He has a strong drive for competence. He seems to have little feel and little sympathy for other people and does not enjoy interacting with others. Self-centered, he nonetheless has a deep moral sense.

Then, the subjects were asked to give a probability that each description belonged to an engineer rather than a lawyer. One group of subjects was told this about the population:

- Condition A: The population consisted of 70 engineers and 30 lawyers.

The second group of subjects was told this about the population:

- Condition B: The population consisted of 30 engineers and 70 lawyers.

Statistically, the probability that a particular description belongs to an engineer rather than a lawyer should be much higher under Condition A than Condition B. However, subjects judged the assignments to be the same in either case. They were basing their judgments solely on the stereotypical personality characteristics of engineers and lawyers, and ignoring the relative probabilities of the two categories.

Interestingly, when subjects were not given any personality description at all and simply asked for the probability that a random individual was an engineer, they answered correctly: $70 \%$ under Condition A and $30 \%$ under Condition B. But when they were given a neutral personality description, one that didn't trigger either stereotype, they assigned the description to an engineer $50 \%$ of the time under both Conditions $A$ and $B$.

And here's a third experiment. Subjects (college students) were given a survey which included these two questions: "How happy are you with your life in general?" and "How many dates did you have last month?" When asked in this order, there was no correlation between the answers. But when asked in the reverse order--when the survey reminded the subjects of how good (or bad) their love life was before asking them about their life in general--there was a $66 \%$ correlation. ${ }^{43}$

Representativeness also explains the base rate fallacy, where people forget that if a particular characteristic is extremely rare, even an accurate test for that characteristic will show false alarms far more often than it will correctly identify the characteristic. Security people run into this heuristic whenever someone tries to sell such things as face scanning, profiling, or data mining as effective ways to find terrorists.

And lastly, representativeness explains the "law of small numbers," where people assume that long-term probabilities also hold in the short run. This is, of course, not true: if the results of three successive coin flips are tails, the odds of heads on the fourth flip are not more than $50 \%$. The coin is not "due" to flip heads. Yet experiments have demonstrated this fallacy in sports betting again and again. ${ }^{44}$

## Cost Heuristics

Humans have all sorts of pathologies involving costs, and this isn't the place to discuss them all. But there are a few specific heuristics I want to summarize, because if we can't evaluate costs right--either monetary costs or more abstract costs--we're not going to make good security trade-offs.

## Mental Accounting

Mental accounting is the process by which people categorize different costs. ${ }^{45}$ People don't simply think of costs as costs; it's much more complicated than that.

Here are the illogical results of two experiments. ${ }^{46}$
In the first, subjects were asked to answer one of these two questions:

- Trade-off 1: Imagine that you have decided to see a play where the admission is $\$ 10$ per ticket. As you enter the theater you discover that you have lost a $\$ 10$ bill. Would you still pay $\$ 10$ for a ticket to the play?
- Trade-off 2: Imagine that you have decided to see a play where the admission is $\$ 10$ per ticket. As you enter the theater you discover that you have lost the ticket. The seat is not marked and the ticket cannot be recovered. Would you pay $\$ 10$ for another ticket?

The results of the trade-off are exactly the same. In either case, you can either see the play and have $\$ 20$ less in your pocket, or not see the play and have $\$ 10$ less in your pocket. But people don't see these trade-offs as the same. Faced with Trade-off $1,88 \%$ of subjects said they would buy the ticket anyway. But faced with Trade-off 2 , only $46 \%$ said they would buy a second ticket. The researchers concluded that there is some sort of mental accounting going on, and the two different $\$ 10$ expenses are coming out of different mental accounts.

The second experiment was similar. Subjects were asked:

- Imagine that you are about to purchase a jacket for $\$ 125$, and a calculator for $\$ 15$. The calculator salesman informs you that the calculator you wish to buy is on sale for $\$ 10$ at the other branch of the store, located 20 minutes' drive away. Would you make the trip to the other store?
- Imagine that you are about to purchase a jacket for $\$ 15$, and a calculator for $\$ 125$. The calculator salesman informs you that the calculator you wish to buy is on sale for $\$ 120$ at the other branch of the store, located 20 minutes' drive away. Would you make the trip to the other store?

Ignore your amazement at the idea of spending $\$ 125$ on a calculator; it's an old experiment. These two questions are basically the same: would you drive 20 minutes to save $\$ 5$ ? But while $68 \%$ of subjects would make the drive to save $\$ 5$ off the $\$ 15$ calculator, only $29 \%$ would make the drive to save $\$ 5$ off the $\$ 125$ calculator.

There's a lot more to mental accounting. ${ }^{47}$ In one experiment, ${ }^{48}$ subjects were asked to imagine themselves lying on the beach on a hot day and how good a cold bottle of their favorite beer would feel. They were to imagine that a friend with them was going up to make a phone call--this was in 1985, before cell phones--and offered to buy them that favorite brand of beer if they gave the friend the money. What was the most the subject was willing to pay for the beer?

Subjects were divided into two groups. In the first group, the friend offered to buy the beer from a fancy resort hotel. In the second group, the friend offered to buy the beer from a rundown grocery store. From a purely economic viewpoint, that should make no difference. The value of one's favorite brand of beer on a hot summer's day has nothing to do with where it was purchased from. (In economic terms, the consumption experience is the same.) But people were willing to pay $\$ 2.65$ on average for the beer from a fancy resort, but only $\$ 1.50$ on average from the run-down grocery store.

The experimenters concluded that people have reference prices in their heads, and that these prices depend on circumstance. And because the reference price was different in the different scenarios, people were willing to pay different amounts. This leads to sub-optimal results. As Thayer writes, "The thirsty beer-drinker who would pay $\$ 4$ for a beer from a resort but only $\$ 2$ from a grocery store will miss out on some pleasant drinking when faced with a grocery store charging $\$ 2.50$."

Researchers have documented all sorts of mental accounting heuristics. Small costs are often not "booked," so people more easily spend money on things like a morning coffee. This is why advertisers often describe large annual costs as "only a few dollars a day." People segregate frivolous money from serious money, so it's easier for them to spend the $\$ 100$ they won in a football pool than a $\$ 100$ tax refund. And people have different mental budgets. In one experiment that illustrates this, ${ }^{49}$ two groups of subjects were asked if they were willing to buy tickets to a play. The first group was told to imagine that they had spent $\$ 50$ earlier in the week on tickets to a basketball game, while the second group was told to
imagine that they had received a $\$ 50$ parking ticket earlier in the week. Those who had spent $\$ 50$ on the basketball game (out of the same mental budget) were significantly less likely to buy the play tickets than those who spent $\$ 50$ paying a parking ticket (out of a different mental budget).

One interesting mental accounting effect can be seen at race tracks. ${ }^{\frac{50}{} \text { Bettors tend to shift }}$ their bets away from favorites and towards long shots at the end of the day. This has been explained by the fact that the average bettor is behind by the end of the day--pari-mutuel betting means that the average bet is a loss--and a long shot can put a bettor ahead for the day. There's a "day's bets" mental account, and bettors don't want to close it in the red.

The effect of mental accounting on security trade-offs isn't clear, but I'm certain we have a mental account for "safety" or "security," and that money spent from that account feels different than money spent from another account. I'll even wager we have a similar mental accounting model for non-fungible costs such as risk: risks from one account don't compare easily with risks from another. That is, we are willing to accept considerable risks in our leisure account--skydiving, knife juggling, whatever--when we wouldn't even consider them if they were charged against a different account.

## Time Discounting

"Time discounting" is the term used to describe the human tendency to discount future costs and benefits. It makes economic sense; a cost paid in a year is not the same as a cost paid today, because that money could be invested and earn interest during the year. Similarly, a benefit accrued in a year is worth less than a benefit accrued today.

Way back in 1937, economist Paul Samuelson proposed a discounted-utility model to explain this all. Basically, something is worth more today than it is in the future. It's worth more to you to have a house today than it is to get it in ten years, because you'll have ten more years' enjoyment of the house. Money is worth more today than it is years from now; that's why a bank is willing to pay you to store it with them.

The discounted utility model assumes that things are discounted according to some rate. There's a mathematical formula for calculating which is worth more--\$100 today or \$120 in twelve months--based on interest rates. Today, for example, the discount rate is $6.25 \%$, meaning that $\$ 100$ today is worth the same as $\$ 106.25$ in twelve months. But of course, people are much more complicated than that.

There is, for example, a magnitude effect: smaller amounts are discounted more than larger ones. In one experiment, ${ }^{51}$ subjects were asked to choose between an amount of money today or a greater amount in a year. The results would make any banker shake his head in wonder. People didn't care whether they received $\$ 15$ today or $\$ 60$ in twelve months. At the same time, they were indifferent to receiving $\$ 250$ today or $\$ 350$ in twelve months, and $\$ 3,000$ today or $\$ 4,000$ in twelve months. If you do the math, that implies a discount rate of
$139 \%, 34 \%$, and $29 \%$--all held simultaneously by subjects, depending on the initial dollar amount.

This holds true for losses as well, ${ }^{\underline{52}}$ although gains are discounted more than losses. In other words, someone might be indifferent to $\$ 250$ today or $\$ 350$ in twelve months, but would much prefer a $\$ 250$ penalty today to a $\$ 350$ penalty in twelve months. Notice how time discounting interacts with prospect theory here.

Also, preferences between different delayed rewards can flip, depending on the time between the decision and the two rewards. Someone might prefer $\$ 100$ today to $\$ 110$ tomorrow, but also prefer $\$ 110$ in 31 days to $\$ 100$ in thirty days.

Framing effects show up in time discounting, too. You can frame something either as an acceleration or a delay from a base reference point, and that makes a big difference. In one experiment, ${ }^{53}$ subjects who expected to receive a VCR in twelve months would pay an average of $\$ 54$ to receive it immediately, but subjects who expected to receive the VCR immediately demanded an average \$126 discount to delay receipt for a year. This holds true for losses as well: people demand more to expedite payments than they would pay to delay them. ${ }^{54}$

Reading through the literature, it sometimes seems that discounted utility theory is full of nuances, complications, and contradictions. Time discounting is more pronounced in young people, people who are in emotional states--fear is certainly an example of this--and people who are distracted. But clearly there is some mental discounting going on; it's just not anywhere near linear, and not easily formularized.

## Heuristics that Affect Decisions

And finally, there are biases and heuristics that affect trade-offs. Like many other heuristics we've discussed, they're general, and not specific to security. But they're still important.

First, some more framing effects.
Most of us have anecdotes about what psychologists call the "context effect": preferences among a set of options depend on what other options are in the set. This has been confirmed in all sorts of experiments--remember the experiment about what people were willing to pay for a cold beer on a hot beach--and most of us have anecdotal confirmation of this heuristic.

For example, people have a tendency to choose options that dominate other options, or compromise options that lie between other options. If you want your boss to approve your \$1M security budget, you'll have a much better chance of getting that approval if you give him a choice among three security plans--with budgets of $\$ 500 \mathrm{~K}, \$ 1 \mathrm{M}$, and $\$ 2 \mathrm{M}$, respectively--than you will if you give him a choice among three plans with budgets of \$250K, \$500K, and \$1M.

The rule of thumb makes sense: avoid extremes. It fails, however, when there's an intelligence on the other end, manipulating the set of choices so that a particular one doesn't seem extreme.
"Choice bracketing" is another common heuristic. In other words: choose a variety. Basically, people tend to choose a more diverse set of goods when the decision is bracketed more broadly than they do when it is bracketed more narrowly. For example, ${ }^{55}$ in one experiment students were asked to choose among one of six different snacks that they would receive at the beginning of the next three weekly classes. One group had to choose the three weekly snacks in advance, while the other group chose at the beginning of each class session. Of the group that chose in advance, $64 \%$ chose a different snack each week, but only $9 \%$ of the group that chose each week did the same.

The narrow interpretation of this experiment is that we overestimate the value of variety. Looking ahead three weeks, a variety of snacks seems like a good idea, but when we get to the actual time to enjoy those snacks, we choose the snack we like. But there's a broader interpretation as well, one borne out by similar experiments and directly applicable to risk taking: when faced with repeated risk decisions, evaluating them as a group makes them feel less risky than evaluating them one at a time. Back to finance, someone who rejects a particular gamble as being too risky might accept multiple identical gambles.

Again, the results of a trade-off depend on the context of the trade-off.
It gets even weirder. Psychologists have identified an "anchoring effect," whereby decisions are affected by random information cognitively nearby. In one experiment ${ }^{56}$, subjects were shown the spin of a wheel whose numbers ranged from 0 and 100, and asked to guess whether the number of African nations in the UN was greater or less than that randomly generated number. Then, they were asked to guess the exact number of African nations in the UN.

Even though the spin of the wheel was random, and the subjects knew it, their final guess was strongly influenced by it. That is, subjects who happened to spin a higher random number guessed higher than subjects with a lower random number.

Psychologists have theorized that the subjects anchored on the number in front of them, mentally adjusting it for what they thought was true. Of course, because this was just a guess, many people didn't adjust sufficiently. As strange as it might seem, other experiments have confirmed this effect.

And if you're not completely despairing yet, here's another experiment that will push you over the edge. ${ }^{57}$ In it, subjects were asked one of these two questions:

- Question 1: Should divorce in this country be easier to obtain, more difficult to obtain, or stay as it is now?
- Question 2: Should divorce in this country be easier to obtain, stay as it is now, or be more difficult to obtain?

In response to the first question, $23 \%$ of the subjects chose easier divorce laws, $36 \%$ chose more difficult divorce laws, and $41 \%$ said that the status quo was fine. In response to the second question, $26 \%$ chose easier divorce laws, $46 \%$ chose more difficult divorce laws, and $29 \%$ chose the status quo. Yes, the order in which the alternatives are listed affects the results.

There are lots of results along these lines, including the order of candidates on a ballot.
Another heuristic that affects security trade-offs is the "confirmation bias." People are more likely to notice evidence that supports a previously held position than evidence that discredits it. Even worse, people who support position A sometimes mistakenly believe that anti-A evidence actually supports that position. There are a lot of experiments that confirm this basic bias and explore its complexities.

If there's one moral here, it's that individual preferences are not based on predefined models that can be cleanly represented in the sort of indifference curves you read about in microeconomics textbooks; but instead, are poorly defined, highly malleable, and strongly dependent on the context in which they are elicited. Heuristics and biases matter. A lot.

This all relates to security because it demonstrates that we are not adept at making rational security trade-offs, especially in the context of a lot of ancillary information designed to persuade us one way or another.

## Making Sense of the Perception of Security

We started out by teasing apart the security trade-off, and listing five areas where perception can diverge from reality:

1. The severity of the risk.
2. The probability of the risk.
3. The magnitude of the costs.
4. How effective the countermeasure is at mitigating the risk.
5. The trade-off itself.

Sometimes in all the areas, and all the time in area 4, we can explain this divergence as a consequence of not having enough information. But sometimes we have all the information and still make bad security trade-offs. My aim was to give you a glimpse of the complicated brain systems that make these trade-offs, and how they can go wrong.

Of course, we can make bad trade-offs in anything: predicting what snack we'd prefer next week or not being willing to pay enough for a beer on a hot day. But security trade-offs are particularly vulnerable to these biases because they are so critical to our survival. Long
before our evolutionary ancestors had the brain capacity to consider future snack preferences or a fair price for a cold beer, they were dodging predators and forging social ties with others of their species. Our brain heuristics for dealing with security are old and well-worn, and our amygdalas are even older.

What's new from an evolutionary perspective is large-scale human society, and the new security trade-offs that come with it. In the past I have singled out technology and the media as two aspects of modern society that make it particularly difficult to make good security trade-offs--technology by hiding detailed complexity so that we don't have the right information about risks, and the media by producing such available, vivid, and salient sensory input--but the issue is really broader than that. The neocortex, the part of our brain that has to make security trade-offs, is, in the words of Daniel Gilbert, "still in beta testing."

I have just started exploring the relevant literature in behavioral economics, the psychology of decision making, the psychology of risk, and neuroscience. Undoubtedly there is a lot of research out there for me still to discover, and more fascinatingly counterintuitive experiments that illuminate our brain heuristics and biases. But already I understand much more clearly why we get security trade-offs so wrong so often.

When I started reading about the psychology of security, I quickly realized that this research can be used both for good and for evil. The good way to use this research is to figure out how humans' feelings of security can better match the reality of security. In other words, how do we get people to recognize that they need to question their default behavior? Giving them more information seems not to be the answer; we're already drowning in information, and these heuristics are not based on a lack of information. Perhaps by understanding how our brains processes risk, and the heuristics and biases we use to think about security, we can learn how to override our natural tendencies and make better security trade-offs. Perhaps we can learn how not to be taken in by security theater, and how to convince others not to be taken in by the same.

The evil way is to focus on the feeling of security at the expense of the reality. In his book Influence, ${ }^{58}$ Robert Cialdini makes the point that people can't analyze every decision fully; it's just not possible: people need heuristics to get through life. Cialdini discusses how to take advantage of that; an unscrupulous person, corporation, or government can similarly take advantage of the heuristics and biases we have about risk and security. Concepts of prospect theory, framing, availability, representativeness, affect, and others are key issues in marketing and politics. They're applied generally, but in today's world they're more and more applied to security. Someone could use this research to simply make people feel more secure, rather than to actually make them more secure.

After all my reading and writing, I believe my good way of using the research is unrealistic, and the evil way is unacceptable. But I also see a third way: integrating the feeling and reality of security.

The feeling and reality of security are different, but they're closely related. We make the best security trade-offs--and by that I mean trade-offs that give us genuine security for a reasonable cost--when our feeling of security matches the reality of security. It's when the two are out of alignment that we get security wrong.

In the past, I've criticized palliative security measures that only make people feel more secure as "security theater." But used correctly, they can be a way of raising our feeling of security to more closely match the reality of security. One example is the tamper-proof packaging that started to appear on over-the-counter drugs in the 1980s, after a few highly publicized random poisonings. As a countermeasure, it didn't make much sense. It's easy to poison many foods and over-the-counter medicines right through the seal--with a syringe, for example--or to open and reseal the package well enough that an unwary consumer won't detect it. But the tamper-resistant packaging brought people's perceptions of the risk more in line with the actual risk: minimal. And for that reason the change was worth it.

Of course, security theater has a cost, just like real security. It can cost money, time, capabilities, freedoms, and so on, and most of the time the costs far outweigh the benefits. And security theater is no substitute for real security. Furthermore, too much security theater will raise people's feeling of security to a level greater than the reality, which is also bad. But used in conjunction with real security, a bit of well-placed security theater might be exactly what we need to both be and feel more secure.

1 Bruce Schneier, Beyond Fear: Thinking Sensibly About Security in an Uncertain World, Springer-Verlag, 2003.
$\underline{2}$ David Ropeik and George Gray, Risk: A Practical Guide for Deciding What's Really Safe and What's Really Dangerous in the World Around You, Houghton Mifflin, 2002.
$\underline{3}$ Barry Glassner, The Culture of Fear: Why Americans are Afraid of the Wrong Things, Basic Books, 1999.

4 Paul Slovic, The Perception of Risk, Earthscan Publications Ltd, 2000.
$\underline{5}$ Daniel Gilbert, "If only gay sex caused global warming," Los Angeles Times, July 2, 2006.
6 Jeffrey Kluger, "How Americans Are Living Dangerously," Time, 26 Nov 2006.
$\underline{7}$ Steven Johnson, Mind Wide Open: Your Brain and the Neuroscience of Everyday Life, Scribner, 2004.

8 Daniel Gilbert, "If only gay sex caused global warming," Los Angeles Times, July 2, 2006.
$\underline{9}$ Donald A. Norman, "Being Analog," http://www.jnd.org/dn.mss/being_analog.html.
Originally published as Chapter 7 of The Invisible Computer, MIT Press, 1998.

10 Daniel Kahneman, "A Perspective on Judgment and Choice," American Psychologist, 2003, 58:9, 697-720.

11 Gerg Gigerenzer, Peter M. Todd, et al., Simple Heuristics that Make us Smart, Oxford University Press, 1999.

12 Daniel Kahneman and Amos Tversky, "Prospect Theory: An Analysis of Decision Under Risk," Econometrica, 1979, 47:263-291.

13 Amos Tversky and Daniel Kahneman, "The Framing of Decisions and the Psychology of Choice," Science, 1981, 211: 453-458.

14 Amos Tversky and Daniel Kahneman, "Evidential Impact of Base Rates," in Daniel Kahneman, Paul Slovic, and Amos Tversky (eds.), Judgment Under Uncertainty: Heuristics and Biases, Cambridge University Press, 1982, pp. 153-160.

15 Daniel J. Kahneman, Jack L. Knetsch, and R.H. Thaler, "Experimental Tests of the Endowment Effect and the Coase Theorem," Journal of Political Economy, 1990, 98: 1325 -1348.

16 Jack L. Knetsch, "Preferences and Nonreversibility of Indifference Curves," Journal of Economic Behavior and Organization, 1992, 17: 131-139.

17 Amos Tversky and Daniel Kahneman, "Advances in Prospect Theory: Cumulative Representation of Subjective Uncertainty," Journal of Risk and Uncertainty, 1992, 5:xx, 297 -323.

18 John Adams, "Cars, Cholera, and Cows: The Management of Risk and Uncertainty," CATO Institute Policy Analysis \#335, 1999.

19 David L. Rosenhan and Samuel Messick, "Affect and Expectation," Journal of Personality and Social Psychology, 1966, 3: 38-44.
$\underline{20}$ Neil D. Weinstein, "Unrealistic Optimism about Future Life Events," Journal of Personality and Social Psychology, 1980, 39: 806-820.

21 D. Kahneman, I. Ritov, and D. Schkade, "Economic preferences or attitude expressions? An analysis of dollar responses to public issues," Journal of Risk and Uncertainty, 1999, 19:220-242.
$\underline{22}$ P. Winkielman, R.B. Zajonc, and N. Schwarz, "Subliminal affective priming attributional interventions," Cognition and Emotion, 1977, 11:4, 433-465.
$\underline{23}$ Daniel Gilbert, "If only gay sex caused global warming," Los Angeles Times, July 2, 2006.
$\underline{24}$ Robyn S. Wilson and Joseph L. Arvai, "When Less is More: How Affect Influences Preferences When Comparing Low-risk and High-risk Options," Journal of Risk Research, 2006, 9:2, 165-178.
$\underline{25}$ J. Cohen, The Privileged Ape: Cultural Capital in the Making of Man, Parthenon Publishing Group, 1989.
$\underline{26}$ Paul Slovic, The Perception of Risk, Earthscan Publications Ltd, 2000.
$\underline{27}$ John Allen Paulos, Innumeracy: Mathematical Illiteracy and Its Consequences, Farrar, Straus, and Giroux, 1988.

28 Amos Tversky and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases," Science, 1974, 185:1124-1130.
$\underline{29}$ Bruce Schneier, Beyond Fear: Thinking Sensibly About Security in an Uncertain World, Springer-Verlag, 2003.

30 Barry Glassner, The Culture of Fear: Why Americans are Afraid of the Wrong Things, Basic Books, 1999.

31 Amos Tversky and Daniel Kahneman, "Availability: A Heuristic for Judging Frequency," Cognitive Psychology, 1973, 5:207-232.
$\underline{32}$ John S. Carroll, "The Effect of Imagining an Event on Expectations for the Event: An Interpretation in Terms of the Availability Heuristic," Journal of Experimental Social Psychology, 1978, 14:88-96.

33 Robert M. Reyes, William C. Thompson, and Gordon H. Bower, "Judgmental Biases Resulting from Differing Availabilities of Arguments," Journal of Personality and Social Psychology, 1980, 39:2-12.

34 S. Jim Sherman, Robert B. Cialdini, Donna F. Schwartzman, and Kim D. Reynolds, "Imagining Can Heighten or Lower the Perceived Likelihood of Contracting a Disease: The Mediating Effect of Ease of Imagery," Personality and Social Psychology Bulletin, 1985, 11:118-127.

35 C. K. Morewedge, D.T. Gilbert, and T.D. Wilson, "The Least Likely of Times: How Memory for Past Events Biases the Prediction of Future Events," Psychological Science, 2005, 16:626-630.

36 Cass R. Sunstein, "Terrorism and Probability Neglect," Journal of Risk and Uncertainty, 2003, 26:121-136.

37 Scott Plous, The Psychology of Judgment and Decision Making, McGraw-Hill, 1993.

38 S.E. Taylor and S.T. Fiske, "Point of View and Perceptions of Causality," Journal of Personality and Social Psychology, 1975, 32: 439-445.

39 Paul Slovic, Baruch Fischhoff, and Sarah Lichtenstein, "Rating the Risks," Environment, 1979, 2: 14-20, 36-39.

40 Amos Tversky and Daniel Kahneman, "Extensional vs Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment," Psychological Review, 1983, 90:??, 293-315.

41 Amos Tversky and Daniel Kahneman, "Judgments of and by Representativeness," in Daniel Kahneman, Paul Slovic, and Amos Tversky (eds.), Judgment Under Uncertainty: Heuristics and Biases, Cambridge University Press, 1982.

42 Daniel Kahneman and Amos Tversky, "On the Psychology of Prediction," Psychological Review, 1973, 80: 237-251.

43 Daniel Kahneman and S. Frederick, "Representativeness Revisited: Attribute Substitution in Intuitive Judgement," in T. Gilovich, D. Griffin, and D. Kahneman (eds.), Heuristics and Biases, Cambridge University Press 2002, pp. 49-81.

44 Thomas Gilovich, Robert Vallone, and Amos Tversky, "The Hot Hand in Basketball: On the Misperception of Random Sequences," Cognitive Psychology, 1985, 17: 295-314.

45 Richard H. Thaler, "Toward a Positive Theory of Consumer Choice," Journal of Economic Behavior and Organization, 1980, 1:39-60.

46 Amos Tversky and Daniel Kahneman, "The Framing of Decisions and the Psychology of Choice," Science, 1981, 211:253:258.

47 Richard Thayer, "Mental Accounting Matters," in Colin F. Camerer, George Loewenstein, and Matthew Rabin, eds., Advances in Behavioral Economics, Princeton University Press, 2004.

48 Richard Thayer, "Mental Accounting and Consumer Choice," Marketing Science, 1985, 4:199-214.

49 Chip Heath and Jack B. Soll, "Mental Accounting and Consumer Decisions," Journal of Consumer Research, 1996, 23:40-52.
$\underline{50}$ Muhtar Ali, "Probability and Utility Estimates for Racetrack Bettors," Journal of Political Economy, 1977, 85:803-815.

51 Richard Thayer, "Some Empirical Evidence on Dynamic Inconsistency," Economics Letters, 1981, 8: 201-207.
$\underline{52}$ George Loewenstein and Drazen Prelec, "Anomalies in Intertemporal Choice: Evidence and Interpretation," Quarterly Journal of Economics, 1992, 573-597.

53 George Loewenstein, "Anticipation and the Valuation of Delayed Consumption," Economy Journal, 1987, 97: 666-684.

54 Uri Benzion, Amnon Rapoport, and Joseph Yagel, "Discount Rates Inferred from Decisions: An Experimental Study," Management Science, 1989, 35:270-284.

55 Itamer Simonson, "The Effect of Purchase Quantity and Timing on Variety-Seeking Behavior," Journal of Marketing Research, 1990, 17:150-162.

56 Amos Tversky and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases," Science, 1974, 185: 1124-1131.

57 Howard Schurman and Stanley Presser, Questions and Answers in Attitude Surveys: Experiments on Wording Form, Wording, and Context, Academic Press, 1981.

58 Robert B. Cialdini, Influence: The Psychology of Persuasion, HarperCollins, 1998. Categories: Psychology of Security
$\leftrightarrows$ The Psychology of Security (Part 1)
What Our Top Spy Doesn't Get: Security and Privacy Aren't Opposites $\rightarrow$

Schneier on Security is a personal website. Opinions expressed are not necessarily those of Resilient, an IBM Company.

