


The Psychology of Security (Part 1)

 schneier.com/essays/archives/2008/01/the_psychology_of_se.html

Introduction

Security is both a feeling and a reality. And they're not the same.

The reality of security is mathematical, based on the probability of different risks and the effectiveness of different countermeasures. We can calculate how secure your home is from burglary, based on such factors as the crime rate in the neighborhood you live in and your door-locking habits. We can calculate how likely it is for you to be murdered, either on the streets by a stranger or in your home by a family member. Or how likely you are to be the victim of identity theft. Given a large enough set of statistics on criminal acts, it's not even hard; insurance companies do it all the time.

We can also calculate how much more secure a burglar alarm will make your home, or how well a credit freeze will protect you from identity theft. Again, given enough data, it's easy.

But security is also a feeling, based not on probabilities and mathematical calculations, but on your psychological reactions to both risks and countermeasures. You might feel terribly afraid of terrorism, or you might feel like it's not something worth worrying about. You might feel safer when you see people taking their shoes off at airport metal detectors, or you might not. You might feel that you're at high risk of burglary, medium risk of murder, and low risk of identity theft. And your neighbor, in the exact same situation, might feel that he's at high risk of identity theft, medium risk of burglary, and low risk of murder.

Or, more generally, you can be secure even though you don't feel secure. And you can feel secure even though you're not. The feeling and reality of security are certainly related to each other, but they're just as certainly not the same as each other. We'd probably be better off if we had two different words for them.

This essay is my initial attempt to explore the feeling of security: where it comes from, how it works, and why it diverges from the reality of security.

Four fields of research--two very closely related--can help illuminate this issue. The first is behavioral economics, sometimes called behavioral finance. Behavioral economics looks at human biases--emotional, social, and cognitive--and how they affect economic decisions. The second is the psychology of decision-making, and more specifically bounded rationality, which examines how we make decisions. Neither is directly related to security, but both look at the concept of risk: behavioral economics more in relation to economic risk, and the psychology of decision-making more generally in terms of security risks. But both fields go a long way to explain the divergence between the feeling

and the reality of security and, more importantly, where that divergence comes from.

There is also direct research into the psychology of risk. Psychologists have studied risk perception, trying to figure out when we exaggerate risks and when we downplay them.

A fourth relevant field of research is neuroscience. The psychology of security is intimately tied to how we think: both intellectually and emotionally. Over the millennia, our brains have developed complex mechanisms to deal with threats. Understanding how our brains work, and how they fail, is critical to understanding the feeling of security.

These fields have a lot to teach practitioners of security, whether they're designers of computer security products or implementers of national security policy. And if this paper seems haphazard, it's because I am just starting to scratch the surface of the enormous body of research that's out there. In some ways I feel like a magpie, and that much of this essay is me saying: "Look at this! Isn't it fascinating? Now look at this other thing! Isn't that amazing, too?" Somewhere amidst all of this, there are threads that tie it together, lessons we can learn (other than "people are weird"), and ways we can design security systems that take the feeling of security into account rather than ignoring it.

The Trade-Off of Security

Security is a trade-off. This is something I have written about extensively, and is a notion critical to understanding the psychology of security. There's no such thing as absolute security, and any gain in security always involves some sort of trade-off.

Security costs money, but it also costs in time, convenience, capabilities, liberties, and so on. Whether it's trading some additional home security against the inconvenience of having to carry a key around in your pocket and stick it into a door every time you want to get into your house, or trading additional security from a particular kind of airplane terrorism against the time and expense of searching every passenger, all security is a trade-off.

I remember in the weeks after 9/11, a reporter asked me: "How can we prevent this from ever happening again?" "That's easy," I said, "simply ground all the aircraft."

It's such a far-fetched trade-off that we as a society will never make it. But in the hours after those terrorist attacks, it's exactly what we did. When we didn't know the magnitude of the attacks or the extent of the plot, grounding every airplane was a perfectly reasonable trade-off to make. And even now, years later, I don't hear anyone second-guessing that decision.

It makes no sense to just look at security in terms of effectiveness. "Is this effective against the threat?" is the wrong question to ask. You need to ask: "Is it a good trade-off?" Bulletproof vests work well, and are very effective at stopping bullets. But for most of us, living in lawful and relatively safe industrialized countries, wearing one is not a good trade-off. The additional security isn't worth it: isn't worth the cost, discomfort, or unfashionableness. Move to another part of the world, and you might make a different trade-off.

We make security trade-offs, large and small, every day. We make them when we decide to lock our doors in the morning, when we choose our driving route, and when we decide whether we're going to pay for something via check, credit card, or cash. They're often not the only factor in a decision, but they're a contributing factor. And most of the time, we don't even realize it. We make security trade-offs intuitively.

These intuitive choices are central to life on this planet. Every living thing makes security trade-offs, mostly as a species--evolving this way instead of that way--but also as individuals. Imagine a rabbit sitting in a field, eating clover. Suddenly, he spies a fox. He's going to make a security trade-off: should I stay or should I flee? The rabbits that are good at making these trade-offs are going to live to reproduce, while the rabbits that are bad at it are either going to get eaten or starve. This means that, as a successful species on the planet, humans should be really good at making security trade-offs.

And yet, at the same time we seem hopelessly bad at it. We get it wrong all the time. We exaggerate some risks while minimizing others. We exaggerate some costs while minimizing others. Even simple trade-offs we get wrong, wrong, wrong--again and again. A Vulcan studying human security behavior would call us completely illogical.

The truth is that we're not bad at making security trade-offs. We are very well adapted to dealing with the security environment endemic to hominids living in small family groups on the highland plains of East Africa. It's just that the environment of New York in 2007 is different from Kenya circa 100,000 BC. And so our feeling of security diverges from the reality of security, and we get things wrong.

There are several specific aspects of the security trade-off that can go wrong. For example:

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. How well disparate risks and costs can be compared.

The more your perception diverges from reality in any of these five aspects, the more your perceived trade-off won't match the actual trade-off. If you think that the risk is greater than it really is, you're going to overspend on mitigating that risk. If you think the risk is real but only affects other people--for whatever reason--you're going to underspend. If you overestimate the costs of a countermeasure, you're less likely to apply it when you should, and if you overestimate how effective a countermeasure is, you're more likely to apply it when you shouldn't. If you incorrectly evaluate the trade-off, you won't accurately balance the costs and benefits.

A lot of this can be chalked up to simple ignorance. If you think the murder rate in your town is one-tenth of what it really is, for example, then you're going to make bad security trade-offs. But I'm more interested in divergences between perception and reality that *can't* be explained that easily. Why is it that, even if someone knows that automobiles kill 40,000 people each year in the U.S. alone, and airplanes kill only hundreds worldwide, he is more afraid of airplanes than automobiles? Why is it

that, when food poisoning kills 5,000 people every year and 9/11 terrorists killed 2,973 people in one non-repeated incident, we are spending tens of billions of dollars per year (not even counting the wars in Iraq and Afghanistan) on terrorism defense while the entire budget for the Food and Drug Administration in 2007 is only \$1.9 billion?

It's my contention that these irrational trade-offs can be explained by psychology. That something inherent in how our brains work makes us more likely to be afraid of flying than of driving, and more likely to want to spend money, time, and other resources mitigating the risks of terrorism than those of food poisoning. And moreover, that these seeming irrationalities have a good evolutionary reason for existing: they've served our species well in the past. Understanding what they are, why they exist, and why they're failing us now is critical to understanding how we make security decisions. It's critical to understanding why, as a successful species on the planet, we make so many bad security trade-offs.

Conventional Wisdom About Risk

Most of the time, when the perception of security doesn't match the reality of security, it's because the perception of the risk doesn't match the reality of the risk. We worry about the wrong things: paying too much attention to minor risks and not enough attention to major ones. We don't correctly assess the magnitude of different risks. A lot of this can be chalked up to bad information or bad mathematics, but there are some general pathologies that come up over and over again.

In *Beyond Fear*, I listed five:

- People exaggerate spectacular but rare risks and downplay common risks.
- People have trouble estimating risks for anything not exactly like their normal situation.
- Personified risks are perceived to be greater than anonymous risks.
- People underestimate risks they willingly take and overestimate risks in situations they can't control.
- Last, people overestimate risks that are being talked about and remain an object of public scrutiny.¹

David Ropeik and George Gray have a longer list in their book *Risk: A Practical Guide for Deciding What's Really Safe and What's Really Dangerous in the World Around You*:

- Most people are more afraid of risks that are new than those they've lived with for a while. In the summer of 1999, New Yorkers were extremely afraid of West Nile virus, a mosquito-borne infection that had never been seen in the United States. By the summer of 2001, though the virus continued to show up and make a few people sick, the fear had abated. The risk was still there, but New Yorkers had lived with it for a while. Their familiarity with it helped them see it differently.
- Most people are less afraid of risks that are natural than those that are human-made. Many people are more afraid of radiation from nuclear waste, or cell phones, than they are of radiation from the sun, a far greater risk.

- Most people are less afraid of a risk they choose to take than of a risk imposed on them. Smokers are less afraid of smoking than they are of asbestos and other indoor air pollution in their workplace, which is something over which they have little choice.
- Most people are less afraid of risks if the risk also confers some benefits they want. People risk injury or death in an earthquake by living in San Francisco or Los Angeles because they like those areas, or they can find work there.
- Most people are more afraid of risks that can kill them in particularly awful ways, like being eaten by a shark, than they are of the risk of dying in less awful ways, like heart disease--the leading killer in America.
- Most people are less afraid of a risk they feel they have some control over, like driving, and more afraid of a risk they don't control, like flying, or sitting in the passenger seat while somebody else drives.
- Most people are less afraid of risks that come from places, people, corporations, or governments they trust, and more afraid if the risk comes from a source they don't trust. Imagine being offered two glasses of clear liquid. You have to drink one. One comes from Oprah Winfrey. The other comes from a chemical company. Most people would choose Oprah's, even though they have no facts at all about what's in either glass.
- We are more afraid of risks that we are more aware of and less afraid of risks that we are less aware of. In the fall of 2001, awareness of terrorism was so high that fear was rampant, while fear of street crime and global climate change and other risks was low, not because those risks were gone, but because awareness was down.
- We are much more afraid of risks when uncertainty is high, and less afraid when we know more, which explains why we meet many new technologies with high initial concern.
- Adults are much more afraid of risks to their children than risks to themselves. Most people are more afraid of asbestos in their kids' school than asbestos in their own workplace.
- You will generally be more afraid of a risk that could directly affect you than a risk that threatens others. U.S. citizens were less afraid of terrorism before September 11, 2001, because up till then the Americans who had been the targets of terrorist attacks were almost always overseas. But suddenly on September 11, the risk became personal. When that happens, fear goes up, even though the statistical reality of the risk may still be very low.²

Others make these and similar points, which are summarized in Table 1.^{3 4 5 6}

When you look over the list in Table 1, the most remarkable thing is how reasonable so many of them seem. This makes sense for two reasons. One, our perceptions of risk are deeply ingrained in our brains, the result of millions of years of evolution. And two, our perceptions of risk are generally pretty good, and are what have kept us alive and reproducing during those millions of years of evolution.

People exaggerate risks that are:

People downplay risks that are:

Table 1: Conventional Wisdom About People and Risk Perception

Spectacular	Pedestrian
-------------	------------

Rare	Common
------	--------

Personified	Anonymous
-------------	-----------

Beyond their control, or externally imposed	More under their control, or taken willingly
---	--

Talked about	Not discussed
--------------	---------------

Intentional or man-made	Natural
-------------------------	---------

Immediate	Long-term or diffuse
-----------	----------------------

Sudden	Evolving slowly over time
--------	---------------------------

Affecting them personally	Affecting others
---------------------------	------------------

New and unfamiliar	Familiar
--------------------	----------

Uncertain	Well understood
-----------	-----------------

Directed against their children	Directed towards themselves
---------------------------------	-----------------------------

Morally offensive	Morally desirable
-------------------	-------------------

Entirely without redeeming features	Associated with some ancillary benefit
-------------------------------------	--

Not like their current situation	Like their current situation
----------------------------------	------------------------------

When our risk perceptions fail today, it's because of new situations that have occurred at a faster rate than evolution: situations that exist in the world of 2007, but didn't in the world of 100,000 BC. Like a squirrel whose predator-evasion techniques fail when confronted with a car, or a passenger pigeon who finds that evolution prepared him to survive the hawk but not the shotgun, our innate capabilities to deal with risk can fail when confronted with such things as modern human society, technology, and

the media. And, even worse, they can be made to fail by others--politicians, marketers, and so on--who exploit our natural failures for their gain.

To understand all of this, we first need to understand the brain.

Risk and the Brain

The human brain is a fascinating organ, but an absolute mess. Because it has evolved over millions of years, there are all sorts of processes jumbled together rather than logically organized. Some of the processes are optimized for only certain kinds of situations, while others don't work as well as they could. And there's some duplication of effort, and even some conflicting brain processes.

Assessing and reacting to risk is one of the most important things a living creature has to deal with, and there's a very primitive part of the brain that has that job. It's the amygdala, and it sits right above the brainstem, in what's called the medial temporal lobe. The amygdala is responsible for processing base emotions that come from sensory inputs, like anger, avoidance, defensiveness, and fear. It's an old part of the brain, and seems to have originated in early fishes. When an animal--lizard, bird, mammal, even you--sees, hears, or feels something that's a potential danger, the amygdala is what reacts immediately. It's what causes adrenaline and other hormones to be pumped into your bloodstream, triggering the fight-or-flight response, causing increased heart rate and beat force, increased muscle tension, and sweaty palms.

This kind of thing works great if you're a lizard or a lion. Fast reaction is what you're looking for; the faster you can notice threats and either run away from them or fight back, the more likely you are to live to reproduce.

But the world is actually more complicated than that. Some scary things are not really as risky as they seem, and others are better handled by staying in the scary situation to set up a more advantageous future response. This means that there's an evolutionary advantage to being able to hold off the reflexive fight-or-flight response while you work out a more sophisticated analysis of the situation and your options for dealing with it.

We humans have a completely different pathway to deal with *analyzing* risk. It's the neocortex, a more advanced part of the brain that developed very recently, evolutionarily speaking, and only appears in mammals. It's intelligent and analytic. It can reason. It can make more nuanced trade-offs. It's also much slower.

So here's the first fundamental problem: we have two systems for reacting to risk--a primitive intuitive system and a more advanced analytic system--and they're operating in parallel. And it's hard for the neocortex to contradict the amygdala.

In his book *Mind Wide Open*, Steven Johnson relates an incident when he and his wife lived in an apartment and a large window blew in during a storm. He was standing right beside it at the time and heard the whistling of the wind just before the window blew. He was lucky--a foot to the side and he

would have been dead--but the sound has never left him:

But ever since that June storm, a new fear has entered the mix for me: the sound of wind whistling through a window. I know now that our window blew in because it had been installed improperly.... I am entirely convinced that the window we have now is installed correctly, and I trust our superintendent when he says that it is designed to withstand hurricane-force winds. In the five years since that June, we have weathered dozens of storms that produced gusts comparable to the one that blew it in, and the window has performed flawlessly.

I know all these facts--and yet when the wind kicks up, and I hear that whistling sound, I can feel my adrenaline levels rise.... Part of my brain--the part that feels most *me*-like, the part that has opinions about the world and decides how to act on those opinions in a rational way--knows that the windows are safe.... But another part of my brain wants to barricade myself in the bathroom all over again.⁷

There's a good reason evolution has wired our brains this way. If you're a higher-order primate living in the jungle and you're attacked by a lion, it makes sense that you develop a lifelong fear of lions, or at least fear lions more than another animal you haven't personally been attacked by. From a risk/reward perspective, it's a good trade-off for the brain to make, and--if you think about it--it's really no different than your body developing antibodies against, say, chicken pox based on a single exposure. In both cases, your body is saying: "This happened once, and therefore it's likely to happen again. And when it does, I'll be ready." In a world where the threats are limited--where there are only a few diseases and predators that happen to affect the small patch of earth occupied by your particular tribe--it works.

Unfortunately, the brain's fear system doesn't scale the same way the body's immune system does. While the body can develop antibodies for hundreds of diseases, and those antibodies can float around in the bloodstream waiting for a second attack by the same disease, it's harder for the brain to

deal with a multitude of lifelong fears.

All this is about the amygdala. The second fundamental problem is that because the analytic system in the neocortex is so new, it still has a lot of rough edges evolutionarily speaking. Psychologist Daniel Gilbert has a great quotation that explains this:

The brain is a beautifully engineered get-out-of-the-way machine that constantly scans the environment for things out of whose way it should right now get. That's what brains did for several hundred million years--and then, just a few million years ago, the mammalian brain learned a new trick: to predict the timing and location of dangers before they actually happened.

Our ability to duck that which is not yet coming is one of the brain's most stunning innovations, and we wouldn't have dental floss or 401(k) plans without it. But this innovation is in the early stages of development. The application that allows us to respond to visible baseballs is ancient and reliable, but the add-on utility that allows us to respond to threats that loom in an unseen future is still in beta testing.⁸

A lot of what I write in the following sections are examples of these newer parts of the brain getting things wrong.

And it's not just risks. People are not computers. We don't evaluate security trade-offs mathematically, by examining the relative probabilities of different events. Instead, we have shortcuts, rules of thumb, stereotypes, and biases--generally known as "heuristics." These heuristics affect how we think about risks, how we evaluate the probability of future events, how we consider costs, and how we make trade-offs. We have ways of generating close-to-optimal answers quickly with limited cognitive capabilities. Don Norman's wonderful essay, "Being Analog," provides a great background for all this.⁹

Daniel Kahneman, who won a Nobel Prize in Economics for some of this work, talks about humans having two separate cognitive systems: one that intuits and one that reasons:

The operations of System 1 are typically fast, automatic, effortless, associative, implicit (not available to introspection), and often emotionally charged; they are also governed by habit and therefore difficult to control or modify. The operations of System 2 are slower, serial, effortful, more likely to be consciously monitored and deliberately controlled; they are also relatively flexible and potentially rule governed.¹⁰

When you read about the heuristics I describe below, you can find evolutionary reasons for why they exist. And most of them are still very useful.¹¹ The problem is that they can fail us, especially in the context of a modern society. Our social and technological evolution has vastly outpaced our evolution as a species, and our brains are stuck with heuristics that are better suited to living in primitive and small family groups.

And when those heuristics fail, our feeling of security diverges from the reality of security.

Risk Heuristics

The first, and most common, area that can cause the feeling of security to diverge from the reality of security is the perception of risk. Security is a trade-off, and if we get the severity of the risk wrong, we're going to get the trade-off wrong. We can do this both ways, of course. We can underestimate some risks, like the risk of automobile accidents. Or we can overestimate some risks, like the risk of a stranger sneaking into our home at night and kidnapping our child. How we get the risk wrong--when we overestimate and when we underestimate--is governed by a few specific brain heuristics.

Prospect Theory

Here's an experiment that illustrates a particular pair of heuristics.¹² Subjects were divided into two groups. One group was given the choice of these two alternatives:

- Alternative A: A sure gain of \$500.
- Alternative B: A 50% chance of gaining \$1,000.

The other group was given the choice of:

- Alternative C: A sure loss of \$500.
- Alternative D: A 50% chance of losing \$1,000.

These two trade-offs aren't the same, but they're very similar. And traditional economics predicts that the difference doesn't make a difference.

Traditional economics is based on something called "utility theory," which predicts that people make trade-offs based on a straightforward calculation of relative gains and losses. Alternatives A and B have the same expected utility: +\$500. And alternatives C and D have the same expected utility: -\$500. Utility theory predicts that people choose alternatives A and C with the same probability and alternatives B and D with the same probability. Basically, some people prefer sure things and others prefer to take chances. The fact that one is gains and the other is losses doesn't affect the mathematics, and therefore shouldn't affect the results.

But experimental results contradict this. When faced with a gain, most people (84%) chose Alternative A (the sure gain) of \$500 over Alternative B (the risky gain). But when faced with a loss, most people (70%) chose Alternative D (the risky loss) over Alternative C (the sure loss).

The authors of this study explained this difference by developing something called "prospect theory." Unlike utility theory, prospect theory recognizes that people have subjective values for gains and losses. In fact, humans have evolved a pair of heuristics that they apply in these sorts of trade-offs. The first is that a sure gain is better than a chance at a greater gain. ("A bird in the hand is better than two in the bush.") And the second is that a sure loss is worse than a chance at a greater loss. Of course, these are not rigid rules--given a choice between a sure \$100 and a 50% chance at \$1,000,000, only a fool would take the \$100--but all things being equal, they do affect how we make trade-offs.

Evolutionarily, presumably it is a better survival strategy to--all other things being equal, of course--accept small gains rather than risking them for larger ones, and risk larger losses rather than accepting smaller losses. Lions chase young or wounded wildebeest because the investment needed to kill them is lower. Mature and healthy prey would probably be more nutritious, but there's a risk of missing lunch entirely if it gets away. And a small meal will tide the lion over until another day. Getting through today is more important than the possibility of having food tomorrow.

Similarly, it is evolutionarily better to risk a larger loss than to accept a smaller loss. Because animals tend to live on the razor's edge between starvation and reproduction, any loss of food--whether small or large--can be equally bad. That is, both can result in death. If that's true, the best option is to risk everything for the chance at no loss at all.

These two heuristics are so powerful that they can lead to logically inconsistent results. Another experiment, the Asian disease problem, illustrates that.¹³ In this experiment, subjects were asked to imagine a disease outbreak that is expected to kill 600 people, and then to choose between two alternative treatment programs. Then, the subjects were divided into two groups. One group was asked to choose between these two programs for the 600 people:

- Program A: "200 people will be saved."
- Program B: "There is a one-third probability that 600 people will be saved, and a two-thirds probability that no people will be saved."

The second group of subjects were asked to choose between these two programs:

- Program C: "400 people will die."
- Program D: "There is a one-third probability that nobody will die, and a two-thirds probability that 600 people will die."

Like the previous experiment, programs A and B have the same expected utility: 200 people saved and 400 dead, A being a sure thing and B being a risk. Same with Programs C and D. But if you read the two pairs of choices carefully, you'll notice that--unlike the previous experiment--they are exactly the same. A equals C, and B equals D. All that's different is that in the first pair they're presented in terms of a gain (lives saved), while in the second pair they're presented in terms of a loss (people dying).

Yet most people (72%) choose A over B, and most people (78%) choose D over C. People make very different trade-offs if something is presented as a gain than if something is presented as a loss.

Behavioral economists and psychologists call this a "framing effect": peoples' choices are affected by how a trade-off is framed. Frame the choice as a gain, and people will tend to be risk averse. But frame the choice as a loss, and people will tend to be risk seeking.

We'll see other framing effects later on.

Another way of explaining these results is that people tend to attach a greater value to changes closer to their current state than they do to changes further away from their current state. Go back to the first pair of trade-offs I discussed. In the first one, a gain from \$0 to \$500 is worth more than a gain from \$500 to \$1,000, so it doesn't make sense to risk the first \$500 for an even chance at a second \$500. Similarly, in the second trade-off, more value is lost from \$0 to -\$500 than from -\$500 to -\$1,000, so it makes sense for someone to accept an even chance at losing \$1,000 in an attempt to avoid losing \$500. Because gains and losses closer to one's current state are worth more than gains and losses further away, people tend to be risk averse when it comes to gains, but risk seeking when it comes to losses.

Of course, our brains don't do the math. Instead, we simply use the mental shortcut.

There are other effects of these heuristics as well. People are not only risk averse when it comes to gains and risk seeking when it comes to losses; people also value something more when it is considered as something that can be lost, as opposed to when it is considered as a potential gain. Generally, the difference is a factor of 2 to 2.5.¹⁴

This is called the "endowment effect," and has been directly demonstrated in many experiments. In one,¹⁵ half of a group of subjects were given a mug. Then, those who got a mug were asked the price at which they were willing to sell it, and those who didn't get a mug were asked what price they were willing to offer for one. Utility theory predicts that both prices will be about the same, but in fact, the median selling price was over twice the median offer.

In another experiment,¹⁶ subjects were given either a pen or a mug with a college logo, both of roughly equal value. (If you read enough of these studies, you'll quickly notice two things. One,

college students are the most common test subject. And two, any necessary props are most commonly purchased from a college bookstore.) Then the subjects were offered the opportunity to exchange the item they received for the other. If the subjects' preferences had nothing to do with the item they received, the fraction of subjects keeping a mug should equal the fraction of subjects exchanging a pen for a mug, and the fraction of subjects keeping a pen should equal the fraction of subjects exchanging a mug for a pen. In fact, most people kept the item they received; only 22% of subjects traded.

And, in general, most people will reject an even-chance gamble (50% of winning, and 50% of losing) unless the possible win is at least twice the size of the possible loss.¹⁷

What does prospect theory mean for security trade-offs? While I haven't found any research that explicitly examines if people make security trade-offs in the same way they make economic trade-offs, it seems reasonable to me that they do at least in part. Given that, prospect theory implies two things. First, it means that people are going to trade off more for security that lets them keep something they've become accustomed to--a lifestyle, a level of security, some functionality in a product or service--than they were willing to risk to get it in the first place. Second, when considering security gains, people are more likely to accept an incremental gain than a chance at a larger gain; but when considering security losses, they're more likely to risk a larger loss than accept the certainty of a small one.

Other Biases that Affect Risk

We have other heuristics and biases about risks. One common one is called "optimism bias": we tend to believe that we'll do better than most others engaged in the same activity. This bias is why we think car accidents happen only to other people, and why we can at the same time engage in risky behavior while driving and yet complain about others doing the same thing. It's why we can ignore network security risks while at the same time reading about other companies that have been breached. It's why we think we can get by where others failed.

Basically, animals have evolved to underestimate loss. Because those who experience the loss tend not to survive, those of us remaining have an evolved experience that losses *don't* happen and that it's okay to take risks. In fact, some have theorized that people have a "risk thermostat," and seek an optimal level of risk regardless of outside circumstances.¹⁸ By that analysis, if something comes along to reduce risk--seat belt laws, for example--people will compensate by driving more recklessly.

And it's not just that we don't think bad things can happen to us, we--all things being equal--believe that good outcomes are more probable than bad outcomes. This bias has been repeatedly illustrated in all sorts of experiments, but I think this one is particularly simple and elegant.¹⁹

Subjects were shown cards, one after another, with either a cartoon happy face or a cartoon frowning face. The cards were random, and the subjects simply had to guess which face was on the next card before it was turned over.

For half the subjects, the deck consisted of 70% happy faces and 30% frowning faces. Subjects faced with this deck were very accurate in guessing the face type; they were correct 68% of the time. The other half was tested with a deck consisting of 30% happy faces and 70% frowning faces. These subjects were much less accurate with their guesses, only predicting the face type 58% of the time. Subjects' preference for happy faces reduced their accuracy.

In a more realistic experiment,²⁰ students at Cook College were asked "Compared to other Cook students--the same sex as you--what do you think are the chances that the following events will happen to you?" They were given a list of 18 positive and 24 negative events, like getting a good job after graduation, developing a drinking problem, and so on. Overall, they considered themselves 15% more likely than others to experience positive events, and 20% less likely than others to experience negative events.

The literature also discusses a "control bias," where people are more likely to accept risks if they feel they have some control over them. To me, this is simply a manifestation of the optimism bias, and not a separate bias.

Another bias is the "affect heuristic," which basically says that an automatic affective valuation--I've seen it called "the emotional core of an attitude"--is the basis for many judgments and behaviors about it. For example, a study of people's reactions to 37 different public causes showed a very strong correlation between 1) the importance of the issues, 2) support for political solutions, 3) the size of the donation that subjects were willing to make, and 4) the moral satisfaction associated with those donations.²¹ The emotional reaction was a good indicator of all of these different decisions.

With regard to security, the affect heuristic says that an overall good feeling toward a situation leads to a lower risk perception, and an overall bad feeling leads to a higher risk perception. This seems to explain why people tend to underestimate risks for actions that also have some ancillary benefit--smoking, skydiving, and such--but also has some weirder effects.

In one experiment,²² subjects were shown either a happy face, a frowning face, or a neutral face, and then a random Chinese ideograph. Subjects tended to prefer ideographs they saw after the happy face, even though the face was flashed for only ten milliseconds and they had no conscious memory of seeing it. That's the affect heuristic in action.

Another bias is that we are especially tuned to risks involving people. Daniel Gilbert again:²³

We are social mammals whose brains are highly specialized for thinking about others. Understanding what others are up to-- what they know and want, what they are doing and planning-- has been so crucial to the survival of our species that our brains have developed an obsession with all things human. We think about people and their intentions; talk about them; look for and remember them.

In one experiment,²⁴ subjects were presented data about different risks occurring in state parks: risks from people, like purse snatching and vandalism, and natural-world risks, like cars hitting deer on the roads. Then, the subjects were asked which risk warranted more attention from state park officials.

Rationally, the risk that causes the most harm warrants the most attention, but people uniformly rated risks from other people as more serious than risks from deer. Even if the data indicated that the risks from deer were greater than the risks from other people, the people-based risks were judged to be more serious. It wasn't until the researchers presented the damage from deer as enormously higher than the risks from other people that subjects decided it deserved more attention.

People are also especially attuned to risks involving their children. This also makes evolutionary sense. There are basically two security strategies life forms have for propagating their genes. The first, and simplest, is to produce a lot of offspring and hope that some of them survive. Lobsters, for example, can lay 10,000 to 20,000 eggs at a time. Only ten to twenty of the hatchlings live to be four weeks old, but that's enough. The other strategy is to produce only a few offspring, and lavish attention on them. That's what humans do, and it's what allows our species to take such a long time to reach maturity. (Lobsters, on the other hand, grow up quickly.) But it also means that we are particularly attuned to threats to our children, children in general, and even other small and cute creatures.²⁵

There is a lot of research on people and their risk biases. Psychologist Paul Slovic seems to have made a career studying them.²⁶ But most of the research is anecdotal, and sometimes the results seem to contradict each other. I would be interested in seeing not only studies about particular heuristics and when they come into play, but how people deal with instances of contradictory heuristics. Also, I would be very interested in research into how these heuristics affect behavior in the context of a strong fear reaction: basically, when these heuristics can override the amygdala and when they can't.

Probability Heuristics

The second area that can contribute to bad security trade-offs is probability. If we get the probability

wrong, we get the trade-off wrong.

Generally, we as a species are not very good at dealing with large numbers. An enormous amount has been written about this, by John Paulos²⁷ and others. The saying goes "1, 2, 3, many," but evolutionarily it makes some amount of sense. Small numbers matter much more than large numbers. Whether there's one mango or ten mangos is an important distinction, but whether there are 1,000 or 5,000 matters less--it's a lot of mangos, either way. The same sort of thing happens with probabilities as well. We're good at 1 in 2 vs. 1 in 4 vs. 1 in 8, but we're much less good at 1 in 10,000 vs. 1 in 100,000. It's the same joke: "half the time, one quarter of the time, one eighth of the time, almost never." And whether whatever you're measuring occurs one time out of ten thousand or one time out of ten million, it's really just the same: almost never.

Additionally, there are heuristics associated with probabilities. These aren't specific to risk, but contribute to bad evaluations of risk. And it turns out that our brains' ability to quickly assess probability runs into all sorts of problems.