

The Psychology of Security—DRAFT

June 28, 2007

Bruce Schneier

13544 words

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:
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

Table 1: Conventional Wisdom About People and Risk Perception

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 a larger loss.

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.

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."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,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.³³ 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 experiment³⁴ divided subjects into two groups, who then read about a fictional disease called “Hyposcemia-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.³⁶ 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.³⁸ 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 is 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 ((reference?)).

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.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 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 from which the descriptions were sampled:

- 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% correlation⁴³

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 run-down 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.⁵⁰ 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,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 \$500K, \$1M, and \$2M, 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 experiment⁵⁶, 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. ((Reference?)) 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 process 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.

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.

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.

⁵ Daniel Gilbert, “If only gay sex caused global warming,” *Los Angeles Times*, 2 Jul 2006.

⁶ Jeffrey Kluger, “How Americans Are Living Dangerously,” *Time*, 26 Nov 2006.

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.

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,” ((citation)).
- ¹⁹ David L. Rosenhan and Samuel Messick, “Affect and Expectation,” *Journal of Personality and Social Psychology*, 1966, 3: 38–44.
- 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.
- ²² P. Winkielman, R.B. Zajonc, and N. Schwarz, “Subliminal affective priming attributional interventions,” *Cognition and Emotion*, 1977, 11:4, 433–465.
- ²³ Daniel Gilbert, “If only gay sex caused global warming,” *Los Angeles Times*, July 2, 2006.
- 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.
- ²⁵ J. Cohen, *The Privileged Ape: Cultural Capital in the Making of Man*, Parthenon Publishing Group, 1989.
- ²⁶ Paul Slovic, *The Perception of Risk*, Earthscan Publications Ltd, 2000.
- 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.
- 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.
- 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, ((volume and page numbers)).
- 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.
- 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.
- 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.