

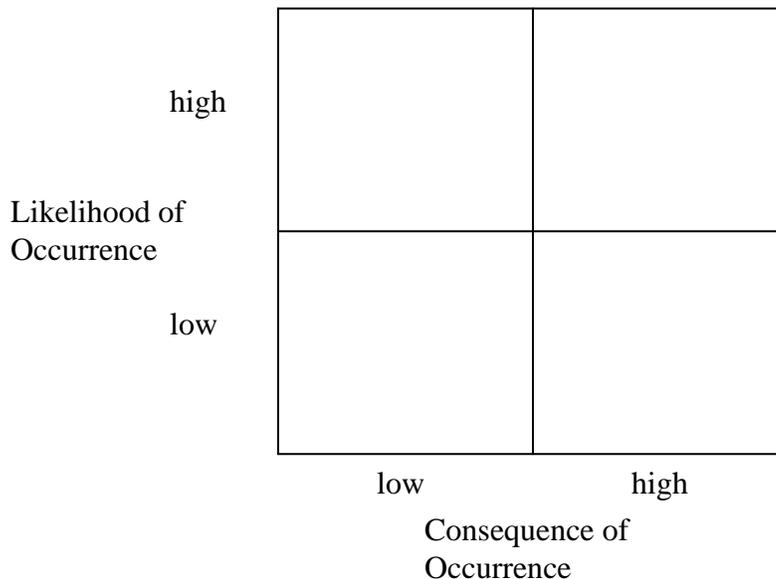
Responsible Citizenship in a Technological Democracy

Note 6: Risk and Risk Perception

Risk

Most people are pretty fuzzy about the notion of risk. In fact, people seem to be especially poor at judging relative risks – they vastly over-react to some risks while simultaneously under-reacting to others. Yet a relatively large number of public policy issues, the kind that a responsible citizen ought to be capable of discussing, involve the notion of risk – what should the government do about mitigating the risk of bird flu, the risk of hurricanes or other natural disasters, the risk of airplane hijacking, the risk of a rogue nation acquiring weapons of mass destruction, and so on. And, in the event that one of these risks actually materializes, what should the government be prepared to do about it. To answer these questions rationally we are going to need to get a better handle on what we mean by the term “risk”, and especially how we compare two or more risks.

The notion of risk is by nature a qualitative and perhaps even subjective (some driving behaviors would be risky for me, but not for a professional NASCAR driver, for example). Nonetheless, experts agree that there are at least two aspects of risk: (1) the likelihood of some event happening, and (2) the consequences of that event. A diagram like the one shown below is often used to illustrate the point (and as part of the explanation for why we are bad at judging relative risks).



As you can see, the diagram is divided into four quadrants. The lower left one corresponds to events that are both rare and inconsequential – they are low risk, no question about it. By contrast, the upper right quadrant corresponds to events that are both quite likely and have dire

consequences – these are *really* risky! The upper left quadrant is for events that are likely, but not especially consequential – pesky you might say. But it’s the lower right quadrant that gives many people the most trouble – unlikely events, but with dire consequences – a large asteroid hitting the earth and causing the extinction of (almost) all life, for example. Talk to your friendly neighborhood dinosaur about that one!

OK, but this diagram is too simplistic for comparing anything but grossly different threats. So, although we know that risk is a qualitative and possibly subjective concept, it is often stated in terms of the following mathematical equation:

Concept: Risk = (Likelihood of Occurrence) × (Consequence of Occurrence)

Yea, Yea, I know that I promised “no equations”, but I figure this one is so simple that it’s under the radar of the equation police. The point is that it defines a numerical value for “risk” and hence allows us to compare the risk of events more finely than the four quadrant diagram. It also has the intuitively comfortable property that if either the Likelihood or the Consequences of an event increase, so does the Risk! But it is a completely arbitrary equation; lots of equations would have this comfortable property. Two other caveats:

1. As we geeks like to say, we’ve just pushed the problem down one level. That is, instead of having to guesstimate the “risk” of an event we now need to guesstimate *two* things – the Likelihood and Consequences of that event. Fortunately both are more easily, and more reliably, estimated. Trust me.
2. If you use the equation, you need to use the same “units” for all the risky events being considered. It doesn’t matter whether you express the Likelihood as a probability¹, as a percentage of the time it happens, or any similar thing. Also, it doesn’t matter whether the Consequences are measured in dollars, or lives lost, or missed opportunities to win the superbowl. What *does* matter is:
 - a. The measures have to increase as either the Likelihood or Consequences increases, and
 - b. You have to use the *same* measure for all risky events being compared.

There are many techniques that collectively go under the name “risk analysis” for determining the Likelihood and Consequences, but we’re not going to go into them here. What is important that you know is that they exist, so that the next time the Transportation Safety Administration (TSA) raises the threat level from orange to red you have every right to ask about the nature of the risk analysis they used. Alas, I’m afraid the TSA method is a wet finger raised into the wind to see which way it’s blowing. They *could* use a sound, scientific method, and perhaps *ought to* – but don’t.

Risk Perception

¹ We’ll talk more about probabilities later.

One reason that TSA doesn't use a more scientific risk analysis is that they are trying to do two things at the same time – one is to reduce real risk and the other is to make passengers *feel* safe. They aren't the same thing. As mentioned in the opening paragraph, people are not good judges of risk, and especially judging the relative risk associated with two or more events. An example that's used all the time to illustrate this is the relative safety of airplanes vs. automobiles. The reality is that, per passenger mile, airplanes are *much* safer than cars, but many people are still afraid of flying, but not of driving.

But there are many, many more examples:

1. Why do people move back to New Orleans when they must know that there will be another flood? Is it, in fact, responsible of the State and Federal governments to support such behavior?
2. In a similar vein, why do people build houses in areas of California that are prone to wind driven grass fires? These fires have happened nearly every year for thousands of years – so they're not exactly a surprise.
3. Why do nearly a third of the population of the US still smoke when the link to lung cancer is so well established?
4. ... and on and on ...

Let's check your perception of risk. Just a few years ago you couldn't turn on a news program, or go to a party for that matter, without hearing about H5N1 Avian Flu (Bird Flu). There was good reason to be concerned. Of the roughly 300 reported cases in humans in the last five years there were about 200 deaths. We don't know how many milder cases were never reported, but among the reported cases, that is a mortality rate of 67%. For calibration, the mortality rate for ordinary seasonal flu is something like one tenth of one percent.

In its current form Avian flu does not easily infect humans and it does not appear to be transmissible from one human to another. But if it ever evolved – and viruses evolve quickly – to be able to infect humans more easily, it would be bad news. So what is riskier Avian flu or ordinary seasonal flu? In particular, if you had fixed resources for research on one or the other, how would you invest them? Think about that a minute before reading on.

OK, the one fact that I didn't give you above is that ordinary flu infects somewhere between 20 and 60 million people each season and kills about 35,000. So my question was asking you to trade off the possibility of an evolved H5N1 that would be horrendous against something that's pesky each year for most of us, but is so pervasive that it is, in fact, a pretty large problem. It's not obvious which to pick – a known killer of 35,000 per year or a killer of 200, but with the possibility of evolving into something much worse. I sure wish I knew more about the likelihood of that evolution!

Over the years quite a number of studies have been done to ascertain why people are poor judges of risk and we can't review them all, but some of the reasons that have been advanced include:

1. *Lack of knowledge*. It seems plausible that experts in a subject have a better basis for judging risks in that subject than a layperson. It's been asserted, for example, that people

who understand nuclear technology favor nuclear power, while those that don't, don't. Studies have shown, however, that while knowledge is a factor, but it isn't the whole story.

2. *The way data is presented.* There are a number of cases of this:
 - a. It seems that if we are told that 95% of patients react favorably to a new drug, people react positively. However, if they are told that 5% of patients have an adverse reaction to the drug, they react negatively. In truth they are both the same data.
 - b. Beware of the use of relative rather than absolute data. If told that a new drug cuts the incidence of death in half, people get really excited. But if told that the new drug cuts the incidence of death from 2% to 1%, it doesn't seem like such a big deal. Again, it may be the same data.
 - c. Beware of data about something related to, but not the *Consequence of concern*. It is often noted that in the Three Mile Island accident, nearly half of the cooling water in the reactor was lost, leading to the meltdown of the reactor core. That sounds really bad. However, loss of coolant and meltdown is not the Consequence of concern – that is, the thing we really ought to be worried about. The Consequence of concern is whether people were exposed to radiation and whether that exposure was likely to affect their health. The answer to that question is “yes”, but not much. About 2 million people received an average dose of about 1 millirem² – which needs to be compared to an average yearly dose of 100-125 millirem from natural radiation from the soil in that part of Pennsylvania.
3. *People tend to be risk averse.* Given two choices for betting on the toss of a coin, one in which you win \$50 if it comes up heads, the other in which you win \$150 if it comes up heads *but* you lose \$100 if it comes up tails, almost everyone will choose the former. Mathematically, the “expected value” is \$25 for both cases.
4. *Cultural associations, biases, and framing.* We all grew up in pretty much the same culture and in that culture some things have negative associations and others have positive ones. If an event becomes associated with one of those negative things, researchers have shown that it will be perceived to be riskier than if it becomes associated with one of the positive ones.
5. *A feeling of familiarity or control.* One reason advanced for the fear of flying, but not driving, is that people feel in control of the car and hence their destiny. Never mind that the pilot and co-pilot are highly trained professionals regularly tested for their proficiency and the typical driver is not.

Aside 1: The “linear, no threshold” (LNT) assumption: Particularly in medicine, but in other contexts as well, we often talk about “dose/response” curves. How much response do we get for a given dose. For example, if an ounce of some toxin will make you sick, generally two ounces will make you even sicker. Or, if one aspirin will make your headache feel somewhat better, two will make it feel even better. Increase the dose and you increase the response.

² Never mind what the unit “millirem” is, except that it's about 1/1000th the dose that would make some people sick.

Well, not quite. Take too much of anything and it's not good for you – that's why the doctor prescribes two aspirin and not 50. Two is enough to relieve most people's headache, but few enough not to have other undesirable effects.

But let's go the other way. If one ounce of a toxin will make you sick, will half an ounce make you half as sick? What about a tenth of an ounce? Or a thousandth of an ounce. Or a billionth? Common sense would say that *really* tiny doses won't have an effect – or, even if they do, that the effect will be so small as to not matter.

Unfortunately, scientists can't agree on where the “threshold” is below which a given dose can be considered harmless. So, in studies of things like toxicity of various substances, they generally make the assumption that the dose/response curve is linear (i.e. a straight line) and that there is no threshold below which the dose can be considered benign. In their heart they know this is later part is wrong, but lacking the ability to say precisely where the threshold is, it seems like the safe assumption to make. This is called the Linear No Threshold, or LNT, assumption.

Unfortunately, every once in a while a well meaning group will fail to understand that the LNT assumption is wrong and will kick up a policy debate. For example, how much arsenic is “safe” in municipal drinking water? Some experts would say 10 parts per billion, others would say 50 – but there will always be a few that say “What? You want to let there be *any* arsenic in *my* drinking water?”

Well, our version of the 2nd law applies here. *Nothing works perfectly. Nothing.* That applies to removing contaminants from drinking water too. Spend enough money and you can make the concentration of contaminants as small as you want, but you can't make it zero. At some point we'll just have to accept that, as best we can tell, concentrations below some level are benign. Maybe someday with better technology we'll have to revise our estimate, either up or down, but as best we can tell today, that's it.

So, beware the LNT assumption. Remember that it's wrong, it is just that we don't know how to find the threshold, so we make this “safe” assumption that there is no threshold.

Aside 2: The cost of risk reduction.

Generally, given the possibility of a risky event, there are things we can do to reduce the risk. We can add anti-lock breaks and airbags to cars, for example, to reduce the risk of injury in an accident. Or, we can tell passengers on the airplane to keep their seatbelts fastened while seated so they don't bang their heads if the plane hits a big air pocket.

It's annoying when the response to a suggestion for reducing some risk is “That's too expensive.” However, although we all feel that our safety is priceless, in fact it isn't. Every risk reduction does have a cost. That cost might be financial, but it could also just be time or convenience, as at the security line at the airport, but it's a cost nonetheless. And in the end, we need to recognize that somebody has to decide how much risk we're willing to tolerate for how much cost.

Our annoyance with the “That’s too expensive” response is that, too often, it’s given without really analyzing the cost or testing whether the market is willing to bear that cost.

Policy Relevance

As noted early in this note, a substantial fraction of public policy issues involve the notion of risk and/or risk perception. Indeed, some of the most important policy issues facing us involve these notions.

There are several things that you ought to remember – all of which are variants of the same basic point, but deserve separate mention.

First, there is an established discipline of risk analysis! More than arm-waving is possible, and you ought to demand it. Get real, quantitative information before you believe much of anything about the riskiness of something.

Second, risk and the perception of risk are not the same thing. Both can be important, however – you have every right to feel safe, and just knowing in your head that you are safe doesn’t always make you feel that way. But, on the other hand, just feeling safe doesn’t mean you are.

Third, be a smart consumer of information about risks. Advocates of various public policies may try to manipulate your perception of certain risks to advance their cause. Some of the techniques to do this are listed above – but don’t treat that list as complete. Creative advocates will always invent new ways to influence you.

Finally, be a bit suspicious of your own perception of a risk – we all suffer from over- or under-reacting to certain risks. Think critically!