From: *The Irrational Economist: Making Decisions in a Dangerous World*, Erwan Michel-Kerjan and Paul Slovic (eds.), New York: Public Affairs Press, 2010, 99-106.

11 Virgin Versus Experienced Risks

CAROLYN KOUSKY, JOHN PRATT, AND RICHARD ZECKHAUSER

INTRODUCTION

Every day we face a multitude of risks. Some we have experienced before; some we have not. Some we have contemplated; others have never crossed our minds. We define four types of risk based on these two dichotomies, as shown in Figure 11.1: virgin risks, contemplated risks, experienced risks, and neglected risks. Individual, institutional, and societal responses to these four types of risk vary in predictable ways.

	Out of Mind	Recognized
No Occurrences	Virgin Risks	Contemplated Risks
Past Occurrences	Neglected Risks	Experienced Risks

FIGURE 11.1 Typology of Risks

Virgin risks are those we have neither experienced nor considered. For an individual, a car crashing into her living room is surely a virgin risk. For society, the meltdown in securities markets of 2008–2009, and its aftermath, was similarly a virgin risk. Of course, if risks are unlikely enough to be virgin risks, we should be concerned with them only if they are of high consequence.

Contemplated risks are those that have not occurred but are recognized. For an individual who has never had a heart attack, the possibility of one may fall into this category. For the United States, the chance of an avian flu pandemic is a contemplated risk: A massive outbreak has not yet happened, but newspapers have discussed this possibility and public and private officials are preparing for it.

Experienced risks, our third category, are risks we both think about and have experienced before. The flu, a fender bender, or a computer crash is an experienced risk for most individuals.

Neglected risks, the last category, have occurred but are not currently contemplated. For example, a century ago an asteroid exploded over Tunguska, Siberia, with 1,000 times the power of the atomic bomb on Hiroshima; it toppled 80 million trees. But the risk of an asteroid explosion has drifted out of public consciousness into the category of neglected risk (leaving aside, of course, NASA's Asteroid and Comet Impact Hazards Group).

This chapter focuses on learning from the occurrence of an extreme event. For many risks, the probability of an occurrence is unknown and may be changing over time. In these cases, even experts must rely on conjectural or subjective assessments of the risk, as we see with climate change. When new information comes to light, such as the occurrence of an extreme event, individuals, institutions, and society must update their assessment of the risk. New information from other sources, such as scientific studies, will also lead to learning. For example, we learned about climate risks from a 2-mile-long ice core from Antarctica that extended the climate record back an additional 210,000 years.

In theory, a rational observer, following the receipt of new information, would update her assessment of the risk given new information. The theoretical model for rational updating, a mathematical formula called Bayes' Rule, is discussed in the next section. Unfortunately, most people are not equipped—mentally or mathematically—to be so ideally rational. Therefore, individuals and society will alter their expectations about extreme events in a nonscientific and often biased fashion. This chapter explores a two-part conjecture: (1) After the occurrence of a virgin risk, people will overestimate the probability of another occurrence in the near future; (2) by contrast, after an experienced risk occurs, people will under-update their assessment of another event occurring soon.

THE INABILITY TO USE BAYESIAN UPDATING IN EVERYDAY PRACTICE

Risks are often posited to have an unknown true probability. The textbook model for how to proceed employs Bayes' Rule (after eighteenth-century British mathematician Thomas Bayes), which shows mathematically how people should rationally change their existing beliefs about something in light of new evidence. Individuals use information available beforehand to form a so-called prior belief about the probability that an event will occur in a given period. New evidence about the risk is captured in something called a likelihood function, which expresses how plausible the evidence is given each possible value of the probability. The prior belief and the likelihood function are then combined to produce what is called a posterior distribution, which is simply the updated version of the prior probability. This requires thinking about a probability model and all its important parameters, and assessing precisely what their distribution was before the occurrence.

While this process might work well in theory, it is rarely used in a complete way in everyday practice, even by those with training in decision analysis. We mention a few prominent explanations among many. Individuals and policy makers are unlikely to have prior beliefs in their heads, especially for virgin risks, and not thinking about every possible risk before it occurs is a perfectly sensible way to organize one's life. However, people often neglect risks that they should consider. For example, in the late 1970s, Howard Kunreuther undertook what would become an important study on flood and earthquake insurance purchases in California. He discovered that people in flood- and earthquake-prone areas often neglected the risk, failing to purchase insurance even when it was subsidized. By doing so, these individuals were ignoring not a 1 in 100,000 possibility but perhaps a 1 in 100 or even a 1 in 50 possibility with grave consequences that could have been alleviated affordably. It is somewhat striking to realize that even today, although information about the risk of earthquakes is well known in California, fewer than 1 in 5 households are insured.

One might think that the prior probability of a virgin risk is zero. Once the event in question occurs, however, it can happen again. A rational person confronted with the occurrence of a virgin-risk event would surely reconsider what his prior assessment of the risk should have been had he thought about the possibility. Although Carl Glatfelter Jr. of McSherrystown, Pennsylvania, never imagined that a car would crash through the front wall of his home, pushing him out a back window,[1] if he'd been forced to place a probability on it he might have given it a 1 in 100 million chance per year. At least until it happened.

It is notoriously difficult, however, to assess what one's prior assessment of the risk would have been had one thought of the event before it actually occurred. Most Americans had never contemplated the possibility of a terrorist attack in this country at the level of 9/11. After it occurred, an attack of that magnitude—or much worse—was not beyond belief, and indeed was thought to have a positive probability of happening again. Thus, most Americans' subjective assessment of terrorism risk should have been far higher after 9/11 than before, despite measures taken to lower such risks after 9/11, since the near-zero prior probability should overwhelm any subsequent response that may have lowered the objective risk. A survey of law school students published in 2003 by economists Kip Viscusi and Richard Zeckhauser in the *Journal of Risk and Uncertainty*, however, found that around 40 percent of respondents believed their personal risk assessment was higher before the attacks than currently.[2] In another study of professional-school students and undergraduate business students in 2005, they showed that over two-thirds of respondents exhibited the same phenomenon.[3] These respondents experienced a recollection bias, whereby after the occurrence of a low-probability event, one thinks that one's prior risk assessment was much higher than it actually was. This could be due to an attempt to reduce cognitive dissonance, for self-justification, or simply to misremembering.

It may also be a variant of hindsight bias, in which knowing the outcome alters an individual's assessment of how likely it was to have occurred. For example, in a 1975 study by psychologist Baruch Fischhoff, who is also a contributor to this book, subjects were given passages to read about the Gurkha raids on the British in the early 1800s. Some were told how the conflict ended, and others were not. When asked what the probability of occurrence of each outcome was, those who knew the outcome gave it a much higher probability. With such "secondary hindsight bias," individuals are unaware that the occurrence of an event influences what they believe *ex post* that they would have estimated *ex ante*. This bias prevents individuals from accurately reconstructing after an event what their prior assessment of the likelihood of that event really was before it happened, making Bayesian updating especially problematical for virgin risks.

What if individuals could accurately assess the prior probability they would have attached to an event had they thought about it? Fully proper prediction of the future risk requires more—namely a probability distribution over a hypothesized true probability p, which is updated once an event happens. Consider the difficulty. An event occurs. An individual concludes that beforehand

she would have thought its probability was $\bar{p}$. Now she has to go one step deeper and assess how likely she would have thought the various values of the true p were that were averaged to produce $\bar{p}$. This strikes us as almost beyond the bounds of human capability, even for those trained in decision analysis, let alone for those who are not.

The individual may also need to define how much new information the event provides. It is often convenient to think of such situations in terms of binomial trials. Although there are some cases in which an event occurs and knowledge of the number of trials in the data can be taken as given, such as an adverse reaction to a drug in a defined group of patients, often it is unclear how many "trials" have occurred. It is difficult to determine how many people have been exposed to a particular chemical, or how many potential terrorist attacks have been prevented, especially if one is not privy to proprietary or classified information. For all these reasons, then, Bayesian updating is unlikely to be used. How, then, do individuals incorporate new information from the occurrence of an event?

UPDATING AS A FUNCTION OF PREVIOUS EXPERIENCE AND PREVIOUS CONTEMPLATION

We conjecture that individuals will excessively update their assessment of virgin risks after one occurs, and fail to (or barely) update their assessment of experienced risks, even when significant updating is warranted. Some experienced risks are well-understood systems with a large amount of frequency data, justifying little updating. For others, however, we have far fewer observations than we think and/or the risk is changing over time. Then, we suspect, individuals update their risk assessments less than they should.

For example, Carolyn Kousky recently examined the changes in property sales prices in St. Louis County, Missouri, before and after the devastating 1993 flood involving the Missouri and Mississippi rivers, both in and out of the floodplain, to determine how homeowners' assessment of the risk changed after the occurrence of an extreme event.[4] During this flood, 100-year floodplains were inundated, as were many 500-year floodplains, particularly those behind failed levees. In our classification, 100-year floodplains represent experienced risks: Floods have previously occurred in these areas and the flood hazard is widely recognized, with national regulations requiring disclosure of the flood risk to homeowners and those with a mortgage required to purchase a flood insurance policy. By contrast, 500-year floodplains represent virgin risks: They had not been recently flooded, and the risk in these areas was also out-of-mind, with no information-disclosure or insurance requirements. After the 1993 flood, property values in the 500-year floodplains declined 2 to 5 percent on average (whether or not they had been flooded in 1993). All property values in municipalities located on one of the flooded rivers declined 6 to 10 percent compared with property in the interior of the county. (These declines were statistically significant.) This finding suggests dramatic updating of the risk. However, there was almost no change in property values in the 100-year floodplains, many of which were flooded, suggesting little to no updating of this experienced risk.

And here is the paradox. Homeowners in the 100-year floodplains should have updated their risk assessments, for three reasons. First, there is insufficient experience to know the probability of a flood with any precision. While there are data on flood events in St. Louis going back to the Lewis and Clark expedition in the early 1800s, this time series is not nearly long enough to provide a tight estimate on the likelihood of an event thought to occur perhaps once a century. Second, the risk is evolving over time. Flood risk has been shown to be increasing over time in the St. Louis area due to structural changes to the Missouri and Mississippi rivers, increased development in the watershed, and possibly climate change. Third, the expected consequence of an event represents a combination of its likelihood and its severity. The floods in 1993 were much more severe than anything previously experienced. There should have been an update in the belief of residents there about the severity of future floods, a factor that would certainly affect housing values.[5]

Two subject areas in behavioral economics, prospect theory and the availability heuristic, help explain the over-updating of virgin risks and the under-updating of experienced risks after an extreme event. A finding of prospect theory is that individuals place excess weight on zero. The Russian Roulette problem illustrates this phenomenon. Most people are willing to pay more to remove one bullet from a six-cylinder gun when it is the only bullet than if there are two (or more) bullets in the gun. That is, a reduction in risk from 1/6 to zero is worth more to them than a reduction from 2/6 to 1/6, even though they are equal reductions in the probability of death, and money is less valuable in the two-bullet case since they are 1/6 likely to die anyway.

Similarly, people perceive an increase in risk from, say, 0 percent to 0.1 percent as large but an equal absolute increase from say 5 percent to 5.1 percent as small. This tendency leads to excessive updating for a previously virgin risk and to barely any updating for an experienced risk. Suppose, for example, an uncontemplated event occurs and fully rational updating would change the risk from 0.01 percent to 1 percent, a 100-fold increase. We conjecture that individuals might instead produce a posterior risk assessment of say 5 percent, a value 5 times too high.

The enormous change in perception when a probability goes from 0 to positive is consistent with evidence from other areas. The theory of just noticeable differences explores such phenomena. For instance, as a noise gets louder, a greater change in volume is needed to make the change perceptible. This is similar to our argument that as base probabilities get larger, small changes in probability are not perceived as well.

The availability heuristic also supports our conjecture. It asserts that individuals assess the probability of an event as higher when examples come to mind more readily. Once an event has occurred, it is much more salient, leading individuals to overestimate its probability. While the first occurrence of a risk makes it suddenly salient, the third occurrence, say, does not add much to its availability. This would explain the substantial updating for virgin risks and the relatively little updating for previously experienced risks.

When dealing with experienced risks, people may suffer from heuristic confusion, assuming, even if incorrect, that they are in a situation where data are extensive and the system is well understood. In such situations, another occurrence of the event in question does not add much information. Many times, however, we act as though we have more information than we do. For instance, in areas where floods occur on average once every 100 years, even if the process were unchanging, thousands of years of data would be needed to accurately assess the probability of a flood in a given year. Such a long time series of data is rarely available.

The bottom line is that with most of the low-probability experienced risks of great interest that affect society as a whole, we have relatively little experience. This means that the updating of our assessments should often be substantial. Indeed, when the probability distribution of a risk is changing over time, the occurrence of an extreme event should lead to even greater updating.

CONCLUSION

We have made two conjectures about human failures when extrapolating from the observance of low-probability, high-consequence events to predictions about future events. First, we tend to overreact when virgin risks occur. The particular danger, now both available and salient, is likely to be overestimated in the future. Second, and by contrast, we tend to raise our probability estimate insufficiently when an experienced risk occurs.

Follow-up research should document these tendencies with many more examples, and in laboratory settings. If improved predictions are our goal, it should also provide rigorous statistical models of effective updating of virgin and experienced risks. Future inquiry should consider resembled risks as well. Evidence from both terrorist incidents and financial markets suggests that we have difficulty extrapolating from risks that, though varied, bear strong similarities.

Behavioral biases such as these are difficult to counteract, but awareness of them is the first step. Requiring careful analysis of all available data could help decision makers to make better risk assessments. History has shown, however, that it is difficult to not overreact to virgin events and to downplay new evidence on experienced risks. Moving from awareness to behavioral change will be challenging, but given the importance of the decisions affected, change is essential.

RECOMMENDED READING

Fischhoff, B. (1975). "Hindsight ≠ Foresight: The Effect of Outcome Knowledge on Judgment Under Uncertainty." *Journal of Experimental Psychology: Human Perception and Performance* 1, no. 3: 288–299.

Kahneman, D., and A. Tversky (1979). "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica* 47, no. 2: 263–292.

Kousky, C. (2010). "Learning from Extreme Events: Risk Perceptions After the Flood." *Land Economics*, forthcoming.

Kunreuther, H. C., ed. (1978). *Disaster Insurance Protection: Public Policy Lessons*. New York: John Wiley and Sons.

Pratt, J. W., and R. J. Zeckhauser (1982). "Inferences from Alarming Events." *Journal of Policy Analysis and Management* 1, no. 3: 371–385.