Harvard Environmental Economics Program

Virgin Versus Experienced Risks

Richard Zeckhauser

Harvard Kennedy School

Carolyn Kousky

Resources for the Future

John Pratt

Harvard Business School

The Harvard Environmental Economics Program

The Harvard Environmental Economics Program develops innovative answers to today's complex environmental issues, by providing a venue to bring together faculty and graduate students from across the University engaged in research, teaching, and outreach in environmental and natural resources economics and related public policy. The program sponsors research projects, develops curricula, and convenes conferences to further understanding of critical issues in environmental and resource economics and policy around the world.

Acknowledgements

Citation Information

Virgin Versus Experienced Risks

Carolyn Kousky, John Pratt, and Richard Zeckhauser

Virgin Versus Experienced Risks

Carolyn Kousky, John Pratt, and Richard Zeckhauser

1. 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 mind. We define four types of risk based on these two dichotomies, as shown in Figure 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.

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 2008-09, and its aftermath, was similarly a virgin risk. Of course, if risks are unlikely enough to be virgin risks, we should only be concerned with them 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 in 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 discuss 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 computer crash is an experienced risk for most individuals.

Neglected risks, the last category, have occurred but are not currently contemplated. For example, exactly 100 years 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).

Figure 1. Typology of Risks

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

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 two-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 centers on a mathematical formula called Bayes' Rule, 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 paper explores a two-part conjecture: (1) after the occurrence of a virgin risk, people will over-estimate 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.

2. 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 what is known as Bayes' Rule (after 18th century British mathematician Thomas Bayes), which shows mathematically how one should 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. These 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 policymakers are unlikely to have priors 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, Kunreuther studied 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 not ignoring 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.

One might think the prior probability of a virgin risk is zero. Once it occurs, however, it can happen again, making the posterior probability positive and thus making a prior probability of zero a logical impossibility. A rational person confronted with the occurrence of a virgin risk 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 had been forced to place a probability on it, he might have given it a 1 in 100 million chance per year.

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'

---

[1] McCune, J. (2008) "Car Crashes Into House, Resident in Serious Condition." The Evening Sun. Hanover, Pennsylvania, September 22.

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 Zeckauser in the Journal of Risk and Uncertainty, however, found around 40% of respondents believed their personal risk assessment was higher before the attacks than currently. 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 (see recommended readings). 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, self-justification, or simply 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 also contributes 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?

3. 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 on 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. During the 1993 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. 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 500-year floodplains declined about 3 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 suggests dramatic updating of the risk. However, there was almost no change in property values in 100-year floodplains, many of which were flooded, suggesting little to no updating of this experienced risk.

And here is the paradox. Homeowners in 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 is 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 certainly should affect housing values.[2]

Two findings from 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 the 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, our argument here is that people perceive an increase in risk from say 0% to 0.1% as large, but an equal absolute increase from say 5% to 5.1% as small. This leads to excessive updating for a previously virgin risk and barely any for an experienced risk. Suppose, for example, an uncontemplated event occurs and fully rational updating would change the risk

---

[2] Note, if the 1993 floods, say, quadrupled expected losses in the two different floodplains, the effect on housing values in the 100-year floodplain should have been much greater, contrary to our findings, but consistent with our conjectures. The factor by which properties in the 500-year floodplain are updated may be greater than in 100-year floodplains, but the absolute amount could be less. For example, houses in the 100-year floodplain may already be reduced in price 18% and those in the 500-year floodplain by 1%. A flood occurs and the homes in the 100-year floodplain have a factor of 3 update, while those in the 500-year floodplain have a factor of 5 update. The factor is thus smaller, but the absolute change in price in larger in the 100-year floodplain.

from 0.01% to 1%, a 100-fold increase. We conjecture that individuals might instead produce a posterior risk assessment of say 5%, 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 things get louder, small changes in volume are harder to perceive. 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 with situations where experience is much more extensive. With extensive data and a well-understood system, another occurrence does not add much information. In many situations, however, we act as though we have more information than we do. For instance, for floods that occur 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 for most of the low probability experienced risks of great interest that affect society as a whole, we have relatively little experience. That means that updating should often be substantial. Moreover, when the probability distribution of a risk is changing over time, the occurrence of an extreme event should lead to even greater updating.

4. Conclusion

We 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, 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 also consider resembled risks. Evidence from both terrorist incidents and financial markets suggests that we have difficulty extrapolating from risks that, though different, bear strong similarities.

Recommended Readings

(1) Fischhoff, B. (1975). "Hindsight ≠ Foresight: The Effect of Outcome Knowledge on Judgment under Uncertainty." Journal of Experimental Psychology: Human Perception and Performance **1**(3): 288-299.

(2) Kahneman, D. and A. Tversky (1979). "Prospect Theory: An Analysis of Decision Under Risk." Econometrica **47**(2): 263-292.

(3) Kousky, C. (2008). Responding to Risk: Information and Decision Making in the Floodplains of St. Louis County, Missouri. Public Policy. Cambridge, MA, Harvard University. **PhD**.

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

(5) Pratt, J. W. and R. J. Zeckhauser (1982). "Inferences from Alarming Events." Journal of Policy Analysis and Management **1**(3): 371-385.

(6) Viscusi, K. W. and R. J. Zeckhauser (2005). "Recollection Bias and the Combat of Terrorism." Journal of Legal Studies **34**(January): 27-55.