ADVANCES IN BEHAVIORAL ECONOMICS[‡]

The Psychology of Tail Events: Progress and Challenges[†]

By NICHOLAS BARBERIS*

Over the past decade, there has been a surge of interest among academic researchers, policymakers, and the broader public in "tail events"—rare, high-impact events. There are many possible reasons for this: the occurrence of particularly salient tail events such as the 9/11 terrorist attacks and the US financial crisis; the rise of the Internet and social media, which have made news and information about tail events much more accessible than before; and the popularization, by Taleb (2007), of the notion of "black swans," to name just a few. In this article, I start by summarizing some recent progress in our understanding of the psychology of tail events. I suggest that much of this progress has centered on the concept of "probability weighting" and, in particular, on applications of this concept in various fields of economics. I then describe some major open questions in this area.

We can think about the psychology of tail events using a two-step framework (Fox and Tversky 2004). In the first step, an individual assesses the probability of a tail event. In the second step, with this probability judgment in hand, he makes a decision. The first step is about beliefs; the second, about preferences. What do we know about how people tackle the task they face at each step?

On the beliefs side, a tentative summary of the available evidence—a summary that papers over many subtleties—is that, when asked to estimate the probability of a tail event, people tend to overestimate this probability. Some of the best evidence on this comes from studies of perceived mortality risks. For example, Lichtenstein et al. (1978) find that people significantly overestimate the frequency of rare causes of death.

On the preferences side, the dominant view is that, if an individual is aware of a potential tail event, he will overweight this potential outcome in his decision making relative to the weight that the outcome would receive in the expected utility framework. This view is most associated with the "probability weighting function," a component of Tversky and Kahneman's (1992) cumulative prospect theory model of decision making under risk. The weighting function transforms subjective probabilities into decision weights, and its principal consequence is to make the individual overweight the tails of any distribution he is considering. The original motivation for this was the simultaneous demand many people have for both lotteries and insurance—they typically prefer a 0.001 chance of $5,000 to a certain $5 but also prefer to pay $5 rather than face a 0.001 chance of a $5,000 loss—a combination of behaviors that is difficult to explain under expected utility. Under probability weighting, however, the unlikely extreme outcomes—gaining or losing $5,000—are overweighted, thereby explaining these choices.[1]

A very rough, first-pass summary of the psychology literature, then, is that a person's

[1] In the original version of prospect theory described in Kahneman and Tversky (1979), probability weighting leads the agent to overweight *all* low-probability outcomes, extreme or not. However, this formulation has significant limitations: it is designed for gambles with at most two non-zero outcomes, and it predicts that people will sometimes choose dominated gambles. Applications of probability weighting in economics have therefore been based almost exclusively on cumulative prospect theory. In this theory, only *extreme* low-probability outcomes are overweighted.

thinking about a tail event is subject to two forces: an event whose true probability, unknown to the individual, is 0.001, say, will first be judged more probable than it actually is—to have probability 0.002, say—and will then be weighted by even more than 0.002 in the individual's decision making: by $\pi(0.002)$, say, where $\pi(0.002) > 0.002$, and where the exact value of $\pi(0.002)$ can be determined from the probability weighting function.

I. Progress

The past few years have seen substantial progress in our understanding of the psychology of tail events. Much of this progress concerns the second of the two steps I described above, in other words, probability weighting. Several papers have used sophisticated methods to carefully estimate the probability weighting function from experimental data; taken together, these studies suggest that, at least in laboratory settings, probability weighting is a remarkably robust feature of risk attitudes (for a review, see Fehr-Duda and Epper 2012). Meanwhile, other papers have linked probability weighting, both theoretically and empirically, to a wide range of economic phenomena. Indeed, in risk-related fields of economics such as finance, insurance, and gambling, there is now more empirical support for probability weighting than for loss aversion, an arguably better-known component of prospect theory.[2] It may be useful to briefly review some of this evidence.[3]

Barberis and Huang (2008) show that, in a financial market where investors evaluate risk according to cumulative prospect theory, probability weighting leads to a new prediction, one that does not emerge from the traditional analysis based on expected utility, namely that the *skewness* in an asset's return distribution will

be priced. For example, a positively skewed stock will be overpriced, relative to the price it would command in an economy with expected utility investors, and will earn a lower average return. The intuition is that, by taking a significant position in a positively skewed stock, an individual gives himself the chance—a small chance, admittedly—of making a lot of money, should the stock experience a right-tail outcome, in other words, should it turn out to be "the next Google." Under probability weighting, this tail event—the event that the stock makes him rich—is overweighted in his decision making. As a result, he is willing to pay a high price for the stock, and to accept a low average return on it.

In the past few years, the prediction that skewness will be priced has received significant empirical support. Several studies, using a variety of techniques for measuring skewness, have found that stocks with high measured skewness do indeed have lower subsequent returns (see, for example, Boyer, Mitton, and Vorkink 2010). Moreover, the idea that skewness is priced has been used to shed light on many puzzling facts: the low average returns of IPO stocks, distressed stocks, bankrupt stocks, stocks traded in over-the-counter markets, and high-volatility stocks (stocks in these categories have positively skewed returns); the apparent overpricing of out-of-the-money options (the returns of these options are highly skewed); the underpricing of conglomerates relative to single-segment firms; and the lack of diversification in many household portfolios.

Probability weighting is also useful for thinking about the *aggregate* stock market—for example, for thinking about the puzzlingly high historical US equity premium. While it has long been known that loss aversion can help resolve this "equity premium puzzle," it has recently become clear that probability weighting can independently generate a large equity premium (De Giorgi and Legg 2012). The reason is that the aggregate stock market is *negatively* skewed: over the course of history, it has been subject to occasional large crashes. Under probability weighting, investors overweight these tail events; as a result, they require a high equity premium.

Insurance markets are also fertile ground for applications of probability weighting. For example, when buying home or automobile

[2] One possible reason for the greater empirical success of probability weighting is that its predictions are relatively robust to the ancillary assumptions that accompany every application of prospect theory—for example, assumptions about the reference point and the degree of narrow framing: an extreme outcome is an extreme outcome, regardless of the specific reference point. By contrast, the predictions of loss aversion can be quite sensitive to the reference point and to the degree of narrow framing.

[3] For more extensive discussion and references regarding the applications mentioned in this section, see Barberis (2013).

insurance, many people select policies with low deductibles, a fact that is hard to make sense of under expected utility. Under probability weighting, however, the unlikely, unpleasant state of the world in which an individual has to make a claim is overweighted in his mind; this pushes him to lower his out-of-pocket expenses in that state, in other words, to choose a lower deductible (Barseghyan et al. forthcoming). Probability weighting can also explain the fact that, while normative models recommend that people allocate a significant fraction of their wealth to annuities upon retirement, few actually do. The idea is that, when thinking about buying an annuity, an individual overweights the unlikely, extreme state of the world in which he dies soon after buying the annuity, thereby receiving much less from the annuity than he initially paid for it. This makes the annuity less attractive (Hu and Scott 2007).

Finally, probability weighting has also proved helpful in understanding betting phenomena such as the favorite-longshot bias, and, more generally, the popularity of casino gambling (Snowberg and Wolfers 2010; Barberis 2012).

In summary, research in psychology suggests that, given some tail event, people overestimate its likelihood; and moreover that, conditional on the probability they assign to the event, people overweight the event in their decision making. This second idea, in the form of probability weighting, has recently emerged as a unifying way of thinking about a broad range of empirical facts. However, notwithstanding this progress, there is still much that we do not know. In the remainder of this article, I discuss some of these open questions.

II. Challenges

Overestimation versus Underestimation.—I noted above that people often overestimate the likelihood of rare, extreme events. Some observers, however, have argued that we sometimes *under*estimate the likelihood of such events. For example, in the run-up to the 2008 US financial crisis, most people appeared to believe that a financial meltdown as severe as the one that eventually occurred was highly improbable. More broadly, the theme of "The Black Swan," Taleb's (2007) much discussed bestseller, is precisely that people tend to assign too low a probability to tail events.

A crucial question, then, is: if we sometimes overestimate probabilities and sometimes underestimate them, does this simply reflect random variation, or are there circumstances in which one mistake is more likely than the other? To tackle this question, it may be helpful to start by thinking about how we come up with likelihood judgments in the first place. One way we do this is through the "availability heuristic" (Tversky and Kahneman 1974): we judge the probability of an event by how easy it is to recall instances of the event. This may explain why we often overestimate the likelihood of tail events: precisely because they are unusual and impactful, these events receive disproportionate media coverage, making it easier for us to recall them and tricking us into judging them more probable than they actually are. Moreover, in many cases, these events generate vivid images and strong emotions—airplane crashes, for example—which keeps them in our memories for longer. However, the availability heuristic can also explain why we sometimes underestimate the likelihood of rare events. For example, many people may have assigned a low probability to a financial crisis in the years leading up to 2008 because it was hard for them to recall a previous instance where economic circumstances like those in 2005 or 2006 were followed in short order by a full-blown crisis.

Overweighting versus Underweighting.—The central implication of probability weighting, implemented through cumulative prospect theory, is that people overweight tail events when they make decisions. However, some researchers have argued that, in some circumstances, people *underweight* tail events. In particular, while agreeing with the widely held view that people overweight tail events when they make decisions "from description," in other words, when they are asked to choose between two *specified* prospects—between a certain $5 and a 0.001 chance of $5,000, say—Hertwig et al. (2004) use laboratory data to argue that people underweight tail events when they make decisions "from experience," in other words, when they are not told the distributions of the gambles they are choosing between but rather have to learn the distributions by sampling from them as often as they like, with replacement. In part, Hertwig et al.'s (2004) effect is driven by sampling error: some participants in their

experiments underweight tail events because, in the small samples they choose to see, the tail event never occurs. However, the effect is not due solely to sampling error, and there is an ongoing effort to understand it more fully.

Earlier, I noted that the overweighting of tail events is consistent with a wide range of empirical facts in finance, insurance, and gambling. How can we reconcile this with Hertwig et al.'s (2004) finding that, in laboratory settings, people underweight tail events when making decisions from experience? There are two possibilities: it may be that, in the field, people are often making decisions from *description*; or it may be that, in the field, people overweight tail events even when they make decisions from experience. In my view, there is evidence for both possibilities. First, it is striking that many of the applications of probability weighting I described above can be thought of as decisions from description, where I include in this category situations in which people use *simple reasoning*, rather than experience, to figure out the distribution of the gamble they are facing. For example, an investor may perceive the returns of distressed stocks as positively skewed not because he has previously experienced or observed these returns, but rather by dint of simple reasoning: "The prices of distressed stocks have been beaten down, so these stocks have limited downside, but a lot of upside—a bit like lottery tickets." Similarly, it is obvious conceptually that the returns of out-of-the-money options are positively skewed; an investor does not need to experience these returns to figure this out.

Second, some of the applications of probability weighting discussed above suggest that, in *field* settings, people overweight tail events even when making decisions from experience. For example, the probability weighting view of the equity premium puzzle is that investors charge a high equity premium because they overweight the occasional stock market crashes they have experienced through time.[4]

Disentangling Beliefs and Preferences.—Earlier, I attributed several financial, insurance, and gambling phenomena to probability weighting, a feature of individual *preferences*. But might these phenomena be driven instead by people's *beliefs*?[5] For example, according to the probability weighting view, investors "overpay" for the typical IPO stock because, while they correctly anticipate the distribution of the stock's future returns, they overweight the state of the world in which the stock turns out to be "the next Google." An alternative view, however, is that investors are over*estimating* the likelihood of the stock being the next Google. Similarly, consumers may choose home insurance policies with low deductibles not because they overweight the state of the world in which they have to make a claim, but because they overestimate the likelihood of this state. Indeed, all of the applications of probability weighting I discussed above can be given a belief-based interpretation—and since people often do overestimate the likelihood of rare events, such an interpretation is not unreasonable. Moreover, the distinction between beliefs and preferences is important: overestimation is a mistake; it is less clear that overweighting is a mistake.

It remains an open question as to whether the phenomena I ascribed above to probability weighting are indeed driven by probability weighting or rather by biased beliefs. To answer this in a direct fashion, we would need reliable data on people's beliefs, and these are not easy to come by. Some indirect evidence, however, favors the preference-based interpretation. Kumar, Page, and Spalt (2011) document that several of the behaviors linked to probability weighting—in particular, the preference for positively skewed assets—are observed more strongly among people who live in Catholic, rather than Protestant, regions of the United States. One interpretation of this finding is that, while people in both Catholic and Protestant regions may initially be drawn to positively skewed assets because of probability weighting, people living in Protestant regions override this initial impulse because buying an asset purely

---

[4] There is another possible type of underweighting that I have not discussed here: if an individual judges a tail event to be sufficiently improbable, he may ignore it altogether in his decision making.

[5] A common misconception is that probability weighting is a feature of beliefs, in other words, that it is about overestimation of small probabilities. It is not; it is a feature of preferences. In the experiments that Kahneman and Tversky (1979) use to motivate probability weighting, participants are told the distributions of the gambles they are facing. Their choices cannot, therefore, be attributed to erroneous beliefs.

for its positive skewness is akin to gambling, a behavior frowned upon by the Protestant church. The Catholic church, however, takes a more lenient view of gambling; this makes it easier for people in Catholic regions to act on their preference for skewness.[6]

Psychological Determinants of Overestimation and Overweighting.—An important challenge that is implicit in the earlier discussion, but that is worth stating explicitly, is that we need a better understanding of *why* people over- or under-estimate the likelihood of tail events; and of *why* they over- or underweight these events in their decision making.

I noted above that, given the central role the availability heuristic plays in judgment, it may be helpful in explaining why people over- or underestimate the likelihood of tail events. Burns, Chiu, and Wu (2011) discuss other possible drivers of these estimation errors, including anchoring and adjustment, and the use of coarse chance categories.

Researchers have also proposed a number of underlying mechanisms for the overweighting of tail events, and, more generally, for the shape of the probability weighting function. For example, Tversky and Kahneman (1992), Rottenstreich and Hsee (2001), and Bordalo, Gennaioli, and Shleifer (2012) discuss mechanisms based on diminishing sensitivity, emotions, and salience, respectively. Nonetheless, it remains an open question as to which of these forces is of primary importance.

Other Challenges.—I end with two more focused research questions. First, there is currently much interest among finance researchers in models where asset prices are determined by investors' time-varying beliefs about the likelihood of an economic disaster. Thus far, the work in this area has not taken account of probability weighting. This is unfortunate: if, as psychology suggests, investors engage in probability weighting, this will significantly affect the quantitative predictions of the "rare disasters" framework.

Second, in dynamic settings, probability weighting generates a time inconsistency: it predicts that how an individual acts in a particular state of the world may differ from how he previously planned to act in that state of the world. While this inconsistency may be useful for understanding real-world behavior—for example, the way people sometimes hold on to losing investments longer than they were planning to—there is relatively little research on it, especially when compared to the large literature on the inconsistency generated by hyperbolic discounting. Barberis (2012) suggests one approach to analyzing probability weighting in dynamic settings, but other approaches are surely also possible.

REFERENCES

Barberis, Nicholas. 2012. "A Model of Casino Gambling." *Management Science* 58 (1): 35–51.

Barberis, Nicholas. 2013. "Thirty Years of Prospect Theory in Economics: A Review and Assessment." *Journal of Economic Perspectives* 27 (1): 173–96.

Barberis, Nicholas, and Ming Huang. 2008. "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices." *American Economic Review* 98 (5): 2066–100.

Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua C. Teitelbaum. Forthcoming. "The Nature of Risk Preferences: Evidence from Insurance Choices." *American Economic Review*.

Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua Teitelbaum. 2013. "Distinguishing Probability Weighting from Risk Perceptions in Field Data." *American Economic Review* 103 (3): TK-TK.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience Theory of Choice under Risk." *Quarterly Journal of Economics* 127 (3): 1243–85.

Boyer, Brian, Todd Mitton, and Keith Vorkink. 2010. "Expected Idiosyncratic Skewness." *Review of Financial Studies* 23 (1): 169–202.

Burns, Zach, Andrew Chiu, and George Wu. 2011. "Overweighting of Small Probabilities." In *Wiley Encyclopedia of Operations Research and Management Science*, edited by James J. Cochran. Hoboken, NJ: Wiley.

De Giorgi, Enrico G., and Shane Legg. 2012. "Dynamic Portfolio Choice and Asset Pricing with Narrow Framing and Probability

---

[6] See Barseghyan et al. (2013) for a recent structural approach to disentangling beliefs and preferences.

Weighting." *Journal of Economic Dynamics and Control* 36 (7): 951–72.

Fehr-Duda, Helga, and Thomas Epper. 2012. "Probability and Risk: Foundations and Economic Implications of Probability-Dependent Risk Preferences." *Annual Review of Economics* 4 (1): 567–93.

Fox, Craig R., and Amos Tversky. 2004. "A Belief-Based Account of Decision under Uncertainty." *Management Science* 44 (7): 879-95.

Hertwig, Ralph, Greg Barron, Elke U. Weber, and Ido Erev. 2004. "Decisions From Experience and the Effect of Rare Events on Risky Choice." *Psychological Science* 15 (8): 534–39.

Hu, Wei-Yin, and Jason S. Scott. 2007. "Behavioral Obstacles in the Annuity Market." *Financial Analysts Journal* 63 (6): 71–82.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47 (2): 263–91.

Kumar, Alok, Jeremy K. Page, and Oliver G. Spalt. 2011. "Religious Beliefs, Gambling Attitudes, and Financial Market Outcomes." *Journal of Financial Economics* 102 (3): 671–708.

Lichtenstein, Sarah, Paul Slovic, Baruch Fischhoff, Mark Layman, and Barbara Combs. 1978. "Judged Frequency of Lethal Events." *Journal of Experimental Psychology and Human Learning* 4 (6): 551-78.

Rottenstreich, Yuval, and Christopher Hsee. 2001. "Money, Kisses, and Electric Shocks: on the Affective Psychology of Risk." *Psychological Science* 12 (3): 185–90.

Snowberg, Erik, and Justin Wolfers. 2010. "Explaining the Favorite-Long Shot Bias: Is It Risk-Love or Misperceptions?" *Journal of Political Economy* 118 (4): 723–46.

Taleb, Nassim Nicholas. 2007. *The Black Swan: The Impact of the Highly Improbable*. New York: Random House.

Tversky, Amos, and Daniel Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185 (4157): 1124–31.

Tversky, Amos, and Daniel Kahneman. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty* 5 (4): 297–323.