

Understanding Uncertainty Shocks and the Role of the Black Swan Frequently Asked Questions

by Anna Orlik and Laura Veldkamp

Q: What is black swan risk and why does it vary over time?

It is a rare event that we may have never before observed. Its estimated probability changes, not because the world changes, but because new data prompts us to re-estimate the distribution that governs the probabilities of all events.

Q: How do we measure black swan risk?

We estimate a smooth distribution that best fits the empirical frequency of observed data. Sometimes called *kernel density estimation*, it is like plotting a histogram of all the GDP data observed in the post-war data sample and then fitting a function to match that histogram.

While kernel density estimation is a fairly standard procedure, we put two twists in that procedure that help us to fit the data and produce forecasts that look more like professional forecasts. First, we estimate non-normal densities that have parameters which regulate skewness. Second, we not only estimate the distribution, but we estimate our uncertainty about the estimates of the distribution. In doing so, we acknowledge that uncertainty comes not only from having volatile outcomes, but also from our own imperfect understanding of the macroeconomy.

Once we have an estimated probability distribution (in fact a family of distributions, each with its own probability), we integrate the probability mass in the left tail of the distribution. That cumulative probability reveals the probability of an extreme negative event.

Q: Why does black swan risk matter?

Volatile and counter-cyclical disaster risk is one of the leading theories of equity risk premia. A vulnerability of this explanation is that we cannot observe the risk and do not have an explanation for why it fluctuates when no disaster is imminent. Estimating black swan risk offers one such explanation. When agents continually re-estimate continuous probability distributions, seemingly innocuous data can change the shape of the estimated distribution in a way that swings tail probabilities around. Using these tools, we can determine whether disaster probabilities fluctuate enough to explain risk premia.

Another important contribution of this set of tools is that they help explain slow economic recoveries in the wake of crises. When agents use data to estimate distributions, new observations create permanent changes in beliefs. A very low growth rate, for example, will never disappear from the data. Forever after observing such an event, agents will continue to have a higher estimate of extreme low growth rates because of its single occurrence. This persistent effect could help resolve one of the main challenges to belief-driven business cycle theories: that shocks fail to have persistent effects. While the economy bounces back quickly from some recessions, it recovers slowly from others because they involve data realizations that permanently change beliefs about the distribution of future outcomes.

Finally, swings in tail probabilities have the largest consequence is for financial claims that are very sensitive to left tail events. Debt is one such asset. Its payoff is constant in most cases, except for default. Thus, the value of debt is sensitive to the risk of an extreme, negative event that would trigger a default. Since most firms have never defaulted, this is a never-before observed event for that firm. Estimating the probability of such an unobserved tail event is exactly what this set of tools is designed to do.

Q: How can I use this tool for my own research project?

The fact that people do not know the true distribution of outcomes has two consequences for probability or uncertainty estimates. The first is that new data will trigger changes in estimates. The second is that these estimates are themselves uncertain. Much of the fluctuation in rare-event risk arises just from re-estimating the distribution, and in particular, adjusting the parameters that govern skewness. Estimating a distribution in real time is simple. Commands like *ksdensity* in MATLAB allow you to do this instantly. A simple approach to adopting this idea is to take the stand that agents act like econometricians and estimate the probability of outcomes each period from realized data. Once they estimate that distribution, they take it as truth until the re-estimate next period. While this myopia, ignoring future re-estimation, creates tension in an otherwise rational model, it has a considerable intellectual history and removes most of the computation complexity from the problem.

Some results, such as the presence of forecast bias, do hinge on having agents who acknowledge their uncertainty about their own estimates. To compute this, we use a standard Metropolis-Hastings algorithm, coupled with a non-linear change of variable at the start. See the appendix of Orlik and Veldkamp (2014) for more details.

[back to research page](#)