

Risk and Resilience For Unknown, Unquantifiable, Systemic, and Unlikely/Catastrophic Threats

Seth D. Baum

Environment, Systems, and Decisions 35(2): 229-236. This version dated 17 June 2015.

Preprint at: http://sethbaum.com/ac/2015_RiskResilience.html.

Background: Risk and Resilience

Risk and resilience are important paradigms for guiding decisions made under uncertainty, in particular decisions about how to protect systems from threats. The risk paradigm tends to emphasize reducing the probabilities and magnitudes of potential losses. The resilience paradigm tends to emphasize increasing the ability of systems to retain critical functionality by absorbing the disturbance, adapting to it, or recovering from it. *This paper discusses the suitability of each paradigm for threats that are unknown, unquantifiable, systemic, and unlikely/catastrophic. The resilience paradigm has sometimes been favored for such threats, but this paper argues that both paradigms are comparably suitable.* The paper uses three examples: Venice during the Black Death plague, superintelligent artificial intelligence (AI), and extraterrestrials (ET) that are much more powerful than humanity.

Unknown Threats

When a threat is completely unknown prior to its occurrence, both risk analysis and resilience analysis are impossible. However, for something to be completely unknown, there must be literally zero available information about it. In practice, it is often possible to identify some information about threats that seem unknown. For example, the Venetians knew there was something spreading a disease, even though they didn't know it was a bacteria. The threats of AI and ET are less well known, but they're also less well suited to the resilience paradigm. If humanity faces an AI or ET that is vastly more powerful, little can be done to increase humanity's resilience.

Unquantifiable Threats

Some threats are known to exist but resist quantification. Their probabilities and/or their magnitudes are deemed unquantifiable. For such threats, calculating risk seems impossible. However, for something to be completely unquantifiable, there must be zero available information about what its quantity might be. In practice, it is often possible to make some quantification, however rough it may be. For example, the Venetians knew the plague was killing many people, even though they lacked modern epidemiology. AI and ET risk is harder to quantify, but still not impossible. For example, AI appears more probable than ET because humans are actively working on AI.

Systemic Threats

Some threats threaten multiple system components or even multiple systems. Sometimes risk analysis only focuses on one component at a time. But this is an error of risk analysis practice, not an error of the risk paradigm itself. Some resilience practice also makes this error. The risk and resilience paradigms are both quite capable of analyzing and managing systemic threats. For example, the Venetians' responded to the plague by quarantining incoming ships. This practice was systemic and can be classified as both risk management and resilience management.

Unlikely/Catastrophic Threats

Some threats are unlikely to occur, but if they do occur, the consequences would be catastrophic. Sometimes risk and resilience analysis neglect these threats. But again, this is an error of practice, not an error of either paradigm. Some risk analysis has been particularly attentive to unlikely/catastrophic risks, including the significant literature on global catastrophic risks. This literature includes some attention to the AI and ET threats, which may be unlikely but would certainly be catastrophic.